



# UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE  
United States Patent and Trademark Office  
Address: COMMISSIONER FOR PATENTS  
P.O. Box 1450  
Alexandria, Virginia 22313-1450  
www.uspto.gov

*Am*

| APPLICATION NO. | FILING DATE | FIRST NAMED INVENTOR | ATTORNEY DOCKET NO. | CONFIRMATION NO. |
|-----------------|-------------|----------------------|---------------------|------------------|
| 09/408,470      | 09/28/1999  | JEFFREY P. KUBALA    | PO9-99-159          | 4085             |

46369 7590 06/06/2005

HESLIN ROTHENBERG FARLEY & MESITI P.C.  
5 COLUMBIA CIRCLE  
ALBANY, NY 12203

EXAMINER

OPIE, GEORGE L

ART UNIT PAPER NUMBER

2194

DATE MAILED: 06/06/2005

Please find below and/or attached an Office communication concerning this application or proceeding.

# Office Action Summary

Application No.

Applicant(s)

09/408,470

Kubala et al.

Examiner

Art Unit

George L. Opie

2126

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --  
Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136 (a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If the period for reply specified above is less than thirty (30) days, a reply within the statutory minimum of thirty (30) days will be considered timely.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133).

## Status

- 1) ☒ Responsive to communication(s) filed on 17 March 2005.
- 2a) ☐ This action is **FINAL**. 2b) ☒ This action is non-final.
- 3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

## Disposition of Claims

- 4) ☒ Claim(s) 1-65 is/are pending in the application.
  - 4a) Of the above claim(s) ☐ is/are withdrawn from consideration.
- 5) ☐ Claim(s) ☐ is/are allowed.
- 6) ☒ Claim(s) 1-65 is/are rejected.
- 7) ☐ Claim(s) ☐ is/are objected to.
- 8) ☐ Claim(s) ☐ are subject to restriction and/or election requirement.

## Application Papers

- 9) ☐ The specification is objected to by the Examiner.
- 10) ☐ The drawing(s) filed on ☐ is/are objected to by the Examiner.
- 11) ☐ The proposed drawing correction filed on ☐ is: a) ☐ approved b) ☐ disapproved.
- 12) ☐ The oath or declaration is objected to by the Examiner.

## Priority under 35 U.S.C. § 119

- 13) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d).
    - a) ☐ All b) ☐ Some \* c) ☐ None of the CERTIFIED copies of the priority documents have been:
      1. ☐ received.
      2. ☐ received in Application No. (Series Code / Serial Number) ☐.
      3. ☐ received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).
- \* See the attached detailed Office action for a list of the certified copies not received.

- 14) ☐ Acknowledgement is made of a claim for domestic priority under 35 U.S.C. & 119(e).

## Attachment(s)

- 14) ☒ Notice of References Cited (PTO-892)
- 15) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
- 16) ☐ Information Disclosure Statement(s) (PTO-1449) Paper No(s) ☐.

- 17) ☐ Interview Summary (PTO-413) Paper No(s) ☐.
- 18) ☐ Notice of Informal Patent Application (PTO-152)
- 19) ☒ Other:  
Text doc for USP6,260,068

10

Art Unit: 2194

## DETAILED ACTION

This Office Action is responsive to the Amendment filed 17 March 2005, in which claims 1, 2, 4-6, 8-13, 16-23, 25-27, 29-34, 37-46, 48-50, 52-57 and 60-65 were amended.

The Office acknowledges Applicant's inclusion of an electronic copy of the amendment on a 3½ inch floppy disk, and the Office would like to thank Applicant for submitting the amendment in electronic form to expedite its processing.

Applicant's Terminal Disclaimer in connection with U.S. Patent 6,587,938 has been received and entered on the instant Application.

1. Applicant should update the related Application information. The cross-referenced Application information must accurately reflect the relevant status of the related cases cited in the Specification. Applicant should review the related applications and provide appropriate amendments to reflect the current information on each referenced patent application.

**2. Request for copy of Applicant's response on floppy disk:**

Please help expedite the prosecution of this application by including, along with your amendment response in paper form, an electronic file copy in WordPerfect, Microsoft Word, or in ASCII text format on a 3½ inch IBM format floppy disk.

Please include all pending claims along with your responsive remarks. Only the paper copy will be entered -- your floppy disk file will be considered a duplicate copy. Signatures are not required on the disk copy. The floppy disk copy is not mandatory; however, it will help expedite the processing of your application. Your cooperation is appreciated.

3. The U.S. Patents used in the art rejections below have been provided as text documents which correspond to the U.S. Patents. The relevant portions of the text documents are cited according to page and line numbers in the art rejections below. For the convenience of Applicant, the cited sections are highlighted in the *text documents*.

4. Claim Rejections - 35 U.S.C. § 103

The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made

Art Unit: 2194

to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

5. Claims 1-5, 22-26, 43 and 45-49 are rejected under 35 U.S.C. § 103(a) as being unpatentable over Zalewski et al. (U.S. Patent 6,260,068).

As to claim 1, Zalewski teaches a method of managing workload of a computing environment (allocation is performed by a software program ... to use the resources appropriately, and provide coordination of resource allocation and sharing, p9 3-20) said method comprising:  
managing workload across two or more partitions of a plurality of partitions of said computing environment (moved from a first partition to a second partition ... in need of a resource, p30 9-41) wherein a partition has one or more central processors allocated thereto (partition must include ... one or more CPUs, p9 27-31)

said managing comprising dynamically redistributing allocation of a shareable resource of at least one partition of said two or more partitions (during runtime an OS instance can temporarily "loan" a CPU to another partition, p21 43-50).

Zalewski does not explicitly state that his migration is conducted for load-balancing purposes, however, Zalewski's teachings clearly suggest that the dynamic resource sharing/migration is performed so as to effectuate achievement of workload goals of the two or more partitions.

As to claim 2, Zalewski teaches the dynamic adjustment is performed transparently to work processing (migration can be initiated and carried out under control of the operating system instances "on the fly" without intervention of the system administrator, p7 20-29).

As to claim 3, Zalewski teaches the shareable resource comprises CPU resources, p21 46-53.

As to claim 4, Zalewski teaches moving at least a portion of the shareable resource from one partition to another partition (moving an I/O processor from one partition to another, p23 19-40).

As to claim 5, Zalewski teaches a "partition priority", p18 19-20 which corresponds to the recitation of managing the resource among two or more partitions based on priority.

As to claims 22-26, note the rejections of claims 1-5 above. Claims 22-26 are the same as claims 1-5, except claims 22-26 are apparatus claims and claims 1-5 are method claims.

Art Unit: 2194

As to claim 43, note the rejection of claim 1 above. The limitations in claim 43 are functionally equivalent to the claim 1 limitations, and would likewise be obvious under the reference and reasoning in the discussion of claim 1 supra.

As to claims 45-49, note the rejections of claims 1-5 above. Claims 45-49 are the same as claims 1-5, except claims 45-49 are apparatus claims and claims 1-5 are method claims.

6. Claims 6-21, 27-33, 34-42, 44 and 50-65 are rejected under 35USC103(a) as being unpatentable over Zalewski as applied to claims 1, 22 and 45 above, and further in view of Maeurer et al. (U.S. Patent 5,301,323).

As to claim 6, Maeurer teaches the resource (channel) "P0 is now shared by CU1 And CU4" p11 3-40, and from this it would have been obvious to provide assigning the shareable resource among two or more partitions based on percentage allocation as recited.

It would have been obvious to combine Maeurer's teachings with Zalewski because the dynamic reallocation/ reconfiguration of a device with respect to channel/partitions provides a flexibility for assigning resources responsive to system load conditions.

As to claim 7, Maeurer (p5 21-54) teaches logical partitions as claimed.

As to claim 8, Maeurer (p3 40-57) teaches adjusting allocation of a plurality of shareable resources (measuring I/O workload and changing the I/O configuration while the system is operating).

As to claim 9, Maeurer (p4 15-32) teaches a dynamic I/O reconfiguration program (DI/OR 34) that corresponds to the workload manager for controlling the dynamic adjustment of resources in the computing environment.

As to claims 10-12, Maeurer teaches the dynamic adjusting/reconfiguring by increasing or decreasing allocation of sharable bandwidth resources, p11 3-40.

As to claim 13, note the claim 1 discussion supra. The limitations in claim 13 are the same as claim 1 but for the fact that claim 13 has an additional limitation of the partitions concurrently share at least one shareable resource. Maeurer's teaching that a resource (channel) "P0 is now shared by CU1 And CU4" p11 3-40 meets the claimed limitation of the partitions concurrently share at least one shareable resource.

As to claim 14, see the claim 3 discussion supra.

Art Unit: 2194

As to claim 15, see Maeurer's p3 40-57 managing/balancing via "workload goals" discussed in the rejection of claim 1 supra.

As to claims 16-17, note the discussions of claims 10-11 above.

As to claims 18-21, see the discussions of claims 4-6 and 8 respectively.

As to claims 27-33, note the rejections of claims 6-12 above. Claims 27-33 are the same as claims 6-12, except claims 27-33 are apparatus claims and claims 6-12 are method claims.

As to claims 34-42, note the rejections of claims 13-21 above. Claims 34-42 are the same as claims 13-21, except claims 34-42 are apparatus claims and claims 13-21 are method claims.

As to claim 44, note the rejection of claim 13 above. The limitations in claim 44 are functionally equivalent to the claim 13 limitations, and would likewise be obvious under the reference and reasoning in the discussion of claim 13 supra.

As to claims 50-56, note the rejections of claims 6-12 above. Claims 50-56 are the same as claims 6-12, except claims 50-56 are computer program product claims and claims 6-12 are method claims.

As to claims 57-65, note the rejections of claims 13-21 above. Claims 57-65 are the same as claims 13-21, except claims 57-65 are computer program product claims and claims 13-21 are method claims.

7. The prior art of record and not relied upon is considered pertinent to the applicant's disclosure. Specifically, the below reference(s) will also have relevancy to one or more elements of the Applicant's claimed invention as follows:

U.S. Patent No. 6,360,303 to Wisler et al. which teaches the partitioning of memory in the sharing of system resources;

U.S. Patent No. 6,173,306 to Raz et al. which teaches the logical mapping of components for optimizing performance responsive to load; and,

U.S. Patent No. 6,101,508 to Wolff which teaches the dynamic allocation -- reassignment of resources for certain processing and maximum scalability.

#### **8. Response to Applicant's Arguments:**

Applicant's remarks accompanying the Amendment filed 17 March 2005, have been considered but are moot in view of the new grounds of rejection.

Art Unit: 2194

**Contact Information:**

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system.

Status information for published applications may be obtained from either Private-PAIR or Public-PAIR.

Status information for unpublished applications is available through Private-PAIR only.

For more information about the PAIR system, see <http://pair-direct.uspto.gov>.

Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

All responses sent by U.S. Mail should be mailed to:

**Commissioner for Patents  
PO Box 1450  
Alexandria, VA 22313-1450**

Hand carried responses should be delivered to the *Customer Service Window* (Randolph Building, 401 Dulany Street, Alexandria, Virginia 22314) and, if submitting an electronic copy on floppy or CD, to expedite its processing, please notify the below identified examiner prior to delivery, so that the Applicant can "handoff" the electronic copy directly to the examiner.

The fax phone number for the organization where this application or proceeding is assigned is 703-872-9306.

Any inquiry of a general nature or relating to the status of this application should be directed to the Group receptionist at **(571) 272-2100**.

Any inquiry concerning this communication or earlier communications from the examiner should be directed to George Opie at (571) 272-3766 or via e-mail at [George.Opie@uspto.gov](mailto:George.Opie@uspto.gov). Internet e-mail should not be used where sensitive data will be exchanged or where there exists a possibility that sensitive data could be identified unless there is an express waiver of the confidentiality requirements under 35 U.S.C. 122 by the Applicant. Sensitive data includes confidential information related to patent applications.

  
**MENG-AL T. AN**  
**SUPERVISORY PATENT EXAMINER**  
**TECHNOLOGY CENTER 2100**

**MAIL WITH OFFICE ACTION**

**ATTACHMENT  
FOR PTO-326**

**U.S. PATENT  
6,260,068**



TITLE: Method and apparatus for migrating resources in a multi-processor computer system

INVENTOR(S): Zalewski, Stephen H., Nashua, NH, United States  
Mason, Andrew H., Hollis, NH, United States  
Jordan, Gregory H., Hollis, NH, United States  
Noel, Karen L., Pembroke, NH, United States  
Kauffman, James R., Nashua, NH, United States

PATENT ASSIGNEE(S): Compaq Computer Corporation, Houston, TX, United States (U.S. corporation)

|                     | NUMBER        | KIND | DATE         |
|---------------------|---------------|------|--------------|
| PATENT INFORMATION: | US 6260068    | B1   | 20010710     |
| APPLICATION INFO.:  | US 1998-95265 |      | 19980610 (9) |
| DOCUMENT TYPE:      | Utility       |      |              |
| FILE SEGMENT:       | GRANTED       |      |              |

|                    | NUMBER     | DATE     | CLASS       | INVENTOR           |
|--------------------|------------|----------|-------------|--------------------|
| REFERENCED PATENT: | US 4843541 | Jun 1989 | 710/036.000 | Bean et al.        |
|                    | US 4853843 | Aug 1989 | 364/200.000 | Ecklund            |
|                    | US 5237566 | Aug 1993 | 370/061.000 | Brand et al.       |
|                    | US 5297262 | Mar 1994 | 710/036.000 | Cox et al.         |
|                    | US 5325517 | Jun 1994 | 395/575.000 | Baker et al.       |
|                    | US 5392397 | Feb 1995 | 709/201.000 | Elko et al.        |
|                    | US 5408649 | Apr 1995 | 395/575.000 | Beshears et al.    |
|                    | US 5414851 | May 1995 | 395/650.000 | Brice, Jr. et al.  |
|                    | US 5450570 | Sep 1995 | 710/010.000 | Richek et al.      |
|                    | US 5471609 | Nov 1995 | 395/575.000 | Yudenfriend et al. |
|                    | US 5481707 | Jan 1996 | 395/650.000 | Murphy, Jr. et al. |
|                    | US 5481719 | Jan 1996 |             | Ackerman et al.    |
|                    | US 5517648 | May 1996 |             | Bertone et al.     |
|                    | US 5537574 | Jul 1996 |             | Elko et al.        |
|                    | US 5574914 | Nov 1996 | 395/650.000 | Hancock et al.     |
|                    | US 5583987 | Dec 1996 |             | Kobayashi et al.   |
|                    | US 5588111 | Dec 1996 |             | Cutts, Jr. et al.  |
|                    | US 5606696 | Feb 1997 |             | Ackerman et al.    |
|                    | US 5613146 | Mar 1997 | 395/800.000 | Gove et al.        |
|                    | US 5625831 | Apr 1997 |             | Priest et al.      |
|                    | US 5636341 | Jun 1997 |             | Matsushita et al.  |
|                    | US 5640584 | Jun 1997 |             | Kandasamy et al.   |
|                    | US 5692193 | Nov 1997 | 709/106.000 | Jagannathan et al. |
|                    | US 5717942 | Feb 1998 |             | Haupt et al.       |
|                    | US 5737763 | Apr 1998 | 711/162.000 | Hilditch           |
|                    | US 5765154 | Jun 1998 | 707/010.000 | Horikiri et al.    |
|                    | US 5784702 | Jul 1998 |             | Greenstein et al.  |
|                    | US 5819020 | Oct 1998 | 714/005.000 | Beeler, Jr.        |
|                    | US 5828894 | Oct 1998 |             | Wilkinson et al.   |
|                    | US 5838968 | Nov 1998 | 709/104.000 | Culbert            |
|                    | US 5860115 | Jan 1999 |             | Neuhard et al.     |
|                    | US 5884018 | Mar 1999 |             | Jardine et al.     |
|                    | US 5898870 | Apr 1999 |             | Okuda et al.       |
|                    | US 5923890 | Jul 1999 |             | Kubala et al.      |
|                    | US 5931938 | Aug 1999 |             | Drogichen et al.   |
|                    | US 5950228 | Sep 1999 |             | Scales et al.      |
|                    | US 5956522 | Sep 1999 |             | Bertone et al.     |
|                    | US 5987621 | Nov 1999 | 714/004.000 | Duso et al.        |
|                    | US 6002851 | Dec 1999 |             | Basavaiah et al.   |
|                    | US 6012151 | Jan 2000 |             | Mano               |

|            |          |                          |
|------------|----------|--------------------------|
| US 6021508 | Feb 2000 | Schmuck et al.           |
| US 6035414 | Mar 2000 | Okazawa et al.           |
| US 6041377 | Mar 2000 | 710/113.000 Mayer et al. |
| US 6047323 | Apr 2000 | 709/227.000 Krause       |
| US 6058423 | May 2000 | 709/226.000 Factor       |
| EP 321694  | Jun 1989 |                          |
| EP 593874  | Apr 1994 |                          |
| WO 8801772 | Mar 1988 |                          |
| WO 9607257 | Mar 1996 |                          |
| WO 9704388 | Feb 1997 |                          |

NON-PATENT REFERENCE:

"DEC's Galaxy Adaptive Partition Multiprocessing,"  
Bannan, Karen, J. PC Week, vol. 14, issue 46, pp.  
61-62, Nov. 1997.\*

G. Hoffnagle, Preface, IBM Systems Journal vol. 36 No.  
2, S/390 Parallel Sysplex Cluster, p. 170, 1997.\*

J.M. Nick, et al., S/390 Cluster Technology: Parallel  
Sysplex, IBM Systems Journal vol. 36, No. 2, S/390  
Parallel Sysplex Cluster, p. 172, 1997.\*

G.M. King et al., Cluster Architectures and S/390  
Parallel Sysplex Scalability IBM Systems Journal vol.  
36, No. 2, S/390 Parallel Sysplex Cluster, p. 221,  
1997.\*

J. Aman, et al., Adaptive Algorithms for Managing a  
Distributed Data Processing Workload, IBM Systems  
Journal vol. 36, No. 2, S/390 Parallel Sysplex Cluster,  
p. 242, 1997.\*

N.S. Bowen, et al, Availability in Parallel Systems:  
Automatic Process Restart, IBM Systems Journal vol. 36,  
No. 2, S/390 Parallel Sysplex Cluster, p. 284, 1997.\*

D. Clitherow, et al., Parallel Sysplex Operational  
Scenarios, IBM Redbooks, 10/98.\*

IBM, S/390 VM/ESA Reference Guide, S/390 VM/ESA  
Reference Guide, 09/98.\*

IBM (website maintained by J.T. Watson Research  
Center), The Hypervisor, IBM Research, Jul. 12, 1996.\*

Sequent's NUMA-Q SMP Architecture, Sequent White Paper,  
1997.\*

J. Chapin, et al., Hive: Fault Containment for  
Shared-Memory Multiprocessors, The 15th ACM Symposium  
on Operating Systems Principles, 12/95.\*

Hive, General Information on OS Homepage.\*

Rohit Chandra, et al., Scheduling and Page Migration  
for Multiprocessor Compute Servers, Sixth International  
Conference on Architectural Support for Programming  
Language and Operating Systems (ASPLOS-VI), 10/1994.

M. Heinrich, et al., The Performance of Flexibility in  
the Stanford FLASH Multiprocessor, Proceedings of the  
ASPLOS-VI, 10/94.

Silicon Graphics, Cellular IRIX 6.4 Technical Report,  
Technical Report, 1996.

Sun Microsystems, Inc., Ultra Enterprise 10000: Dynamic  
System Domains, Technical White Paper, 1997.

Sun Microsystems, Inc., DR Configuration ISSUES,  
Dynamic Reconfiguration User's Guide, 1997.

A. Charlesworth, et al., Gigaplane-XB: Extending the  
Ultra Enterprise Family, Sun Microsystems, Inc., Jul.  
30, 1997.

VAX/VMS Internals and Data Structures, Digital Equipment Corp., 4/81.  
 DEC, TruCluster: Digital's UNIX Cluster, Cluster White Paper, 1996.  
 DEC, Network Multiprocessing Comes of Age.  
 DEC, Symmetrical Multiprocessing, Technical Overview.  
 J.A. Hall, Engineering Excellence: DEC OSF/1 Symmetric Multiprocessing, DEC.  
 Sequent, Implementation and Performance of a CC-Numa System, Sequent White Paper.  
 C.Koppe, NUMA Architectures and User Level Scheduling a Short Introduction, May 14, 1996.  
 C. Hanna, Logical Partitioning Methodologies, CMG (Conference), 1993.  
 D. Bartholomew, The NUMA Invasion, Industry Week, Jan. 6, 1997.  
 Sun Microsystems, Inc., Ultra Enterprise 10000 Server, Technical White Paper, 1997.  
 J. T. Turner, Specials, Skeleton Crew USA, Inc., Jan. 16, 1997.  
 B.F. Smith, DB2 and Business Intelligence in a S/390 Parallel Sysplex, DB2 and Business Intelligence in a S/390 Parallel Sysplex (Query Parallelism on DB2 Version 5), Nov. 11, 1997.  
 Beck, "AAMP: A Multiprocessor Approach for Operating Systems and Application Migration," Operating System Review 24:41-55 (1990).  
 European Search Report dated Jun. 27, 2000 (Application No. 98309009.9-2201-).  
 Nanda et al., "Mapping Applications onto a Cache Coherent Multiprocessor," Proceedings on Supercomputing '92, pp. 368-377 (1992).  
 Rashid et al., "Machine-Independent Virtual Memory Management for Paged Uniprocessor and Multiprocessor Architectures," IEEE Transactions on Computers 37:896-907 (1988).  
 Beck, "AAMP: A Multiprocessor Approach for Operating Systems and Application Migration," Operating System Review 24:41-55 (Apr., 1990).  
 European Search Report dated Jun. 26, 2000 (Application No. 98309011.9-2201-), 4 pages.  
 Ohmori et al., "System management of MICS-II--a virtual machine complex"; Third USA-Japan Computer Conference Proceedings, San Francisco, Calif., Sep. 1978, pp. 425-429.

PRIMARY EXAMINER: Alam, Hosain T.  
 LEGAL REPRESENTATIVE: Williams, Morgan & Amerson  
 NUMBER OF CLAIMS: 42  
 EXEMPLARY CLAIM: 1  
 NUMBER OF DRAWINGS: 14 Drawing Figure(s); 10 Drawing Page(s)  
 ABSTRACT:

Multiple instances of operating systems execute cooperatively in a single multi-processor computer wherein a single physical machine is subdivided by software into multiple partitions, each with the ability to run a distinct copy, or instance, of an operating system. Each individual instance has the resources it needs to execute independently, but instances cooperate to migrate resources from one partition to another. The migration can be initiated and

carried out under control of the operating system instances "on the fly" without intervention of the system administrator. Alternatively, a system administrator can reconfigure the system. Resource migration is carried out under a "push" model in which resources are controlled by an owning partition and must be released by that partition before they can be migrated to another partition. In accordance with this model, a first operating system instance which requires a resource must first request the resource from a second instance. In response to this request, the second instance determines whether it can spare the resource, and if so, begins to bring the resource into an idle state. The resource is actually transferred when the second instance stops using the resource.

#### FIELD OF THE INVENTION

This invention relates to multiprocessor computer architectures in which processors and other computer hardware resources are grouped in partitions, each of which has an operating system instance and, more specifically, to methods and apparatus for migrating computer hardware resources from one partition to another without rebooting the computer system.

#### BACKGROUND OF THE INVENTION

The efficient operation of many applications in present computing environments depend upon fast, powerful and flexible computing systems. The configuration and design of such systems has become very complicated when such systems are to be used in an "enterprise" commercial environment where there may be many separate departments, many different problem types and continually changing computing needs. Users in such environments generally want to be able to quickly and easily change the capacity of the system, its speed and its configuration. They may also want to expand the system work capacity and change configurations to achieve better utilization of resources without stopping execution of application programs on the system. In addition they may want be able to configure the system in order to maximize resource availability so that each application will have an optimum computing configuration.

Traditionally, computing speed has been addressed by using a "shared nothing" computing architecture where data, business logic, and graphic user interfaces are distinct tiers and have specific computing resources dedicated to each tier. Initially, a single central processing unit was used and the power and speed of such a computing system was increased by increasing the clock rate of the single central processing unit. More recently, computing systems have been developed which use several processors working as a team instead one massive processor working alone. In this manner, a complex application can be distributed among many processors instead of waiting to be executed by a single processor. Such systems typically consist of several central processing units (CPUs) which are controlled by a single operating system. In a variant of a multiple processor system called "symmetric multiprocessing" or SMP, the applications are distributed equally across all processors. The processors also share memory. In another variant called "asymmetric multiprocessing" or AMP, one processor acts as a "master" and all of the other processors act as "slaves." Therefore, all operations, including the operating system, must pass through the master before being passed onto the slave processors. These multiprocessing architectures have the advantage that performance can be increased by adding additional processors, but suffer from the disadvantage that the software running on such systems must be carefully written to take advantage of the multiple processors and it is difficult to scale the software as the number of processors increases. Current commercial workloads do not scale well beyond 8-24 CPUs as a single SMP system, the exact number depending upon platform, operating system and application mix.

For increased performance, another typical answer has been to dedicate computer resources (machines) to an application in order to optimally tune the machine resources to the application. However, this approach has not been adopted by

the majority of users because most sites have many applications and separate databases developed by different vendors. Therefore, it is difficult, and expensive, to dedicate resources among all of the applications especially in environments where the application mix is constantly changing. Further, with dedicated resources, it is essentially impossible to quickly and easily migrate resources from one computer system to another, especially if different vendors are involved. Even if such a migration can be performed, it typically involves the intervention of a system administrator and requires at least some of the computer systems to be powered down and rebooted.

Alternatively, a computing system can be partitioned with hardware to make a subset of the resources on a computer available to a specific application. This approach avoids dedicating the resources permanently since the partitions can be changed, but still leaves issues concerning performance improvements by means of load balancing of resources among partitions and resource availability.

The availability and maintainability issues were addressed by a "shared everything" model in which a large centralized robust server that contains most of the resources is networked with and services many small, uncomplicated client network computers. Alternatively, "clusters" are used in which each system or "node" has its own memory and is controlled by its own operating system. The systems interact by sharing disks and passing messages among themselves via some type of communication network. A cluster system has the advantage that additional systems can easily be added to a cluster. However, networks and clusters suffer from a lack of shared memory and from limited interconnect bandwidth which places limitations on performance.

In many enterprise computing environments, it is clear that the two separate computing models must be simultaneously accommodated and each model optimized. Further, it is highly desirable to be able to modify computer configurations "on the fly" without rebooting any of the systems. Several prior art approaches have been used to attempt this accommodation. For example, a design called a "virtual machine" or VM developed and marketed by International Business Machines Corporation, Armonk, N.Y., uses a single physical machine, with one or more physical processors, in combination with software which simulates multiple virtual machines. Each of those virtual machines has, in principle, access to all the physical resources of the underlying real computer. The assignment of resources to each virtual machine is controlled by a program called a "hypervisor". There is only one hypervisor in the system and it is responsible for all the physical resources. Consequently, the hypervisor, not the other operating systems, deals with the allocation of physical hardware. The hypervisor intercepts requests for resources from the other operating systems and deals with the requests in a globally-correct way.

The VM architecture supports the concept of a "logical partition" or LPAR. Each LPAR contains some of the available physical CPUs and resources which are logically assigned to the partition. The same resources can be assigned to more than one partition. LPARs are set up by an administrator statically, but can respond to changes in load dynamically, and without rebooting, in several ways. For example, if two logical partitions, each containing ten CPUs, are shared on a physical system containing ten physical CPUs, and, if the logical ten CPU partitions have complementary peak loads, each partition can take over the entire physical ten CPU system as the workload shifts without a re-boot or operator intervention.

In addition, the CPUs logically assigned to each partition can be turned "on" and "off" dynamically via normal operating system operator commands without re-boot. The only limitation is that the number of CPUs active at system initialization is the maximum number of CPUs that can be turned "on" in any partition.

Finally, in cases where the aggregate workload demand of all partitions is more than can be delivered by the physical system, LPAR "weights" can be used to define the portion of the total CPU resources which is given to each partition. These weights can be changed by system administrators, on-the-fly, with no disruption.

Another prior art system is called a "Parallel Sysplex" and is also marketed and developed by the International Business Machines Corporation. This architecture consists of a set of computers that are clustered via a hardware entity called a "coupling facility" attached to each CPU. The coupling facilities on each node are connected, via a fiber-optic link, and each node operates as a traditional SMP machine, with a maximum of 10 CPUs. Certain CPU instructions directly invoke the coupling facility. For example, a node registers a data structure with the coupling facility, then the coupling facility takes care of keeping the data structures coherent within the local memory of each node.

The Enterprise 10000 Unix server developed and marketed by Sun Microsystems, Mountain View, Calif., uses a partitioning arrangement called "Dynamic System Domains" to logically divide the resources of a single physical server into multiple partitions, or domains, each of which operates as a stand-alone server.

Each of the partitions has CPUs, memory and I/O hardware. Dynamic reconfiguration allows a system administrator to create, resize, or delete domains "on the fly" and without rebooting. Every domain remains logically isolated from any other domain in the system, isolating it completely from any software error or CPU, memory, or I/O error generated by any other domain. There is no sharing of resources between any of the domains.

The Hive Project conducted at Stanford University uses an architecture which is structured as a set of cells. When the system boots, each cell is assigned a range of nodes, each having memory and I/O devices, that the cell owns throughout execution. Each cell manages the processors, memory and I/O devices on those nodes as if it were an independent operating system. The cells cooperate to present the illusion of a single system to user-level processes. Hive cells are not responsible for deciding how to divide their resources between local and remote requests. Each cell is responsible only for maintaining its internal resources and for optimizing performance within the resources it has been allocated. Global resource allocation is carried out by a user-level process called "wax." The Hive system attempts to prevent data corruption by using certain fault containment boundaries between the cells. In order to implement the tight sharing expected from a multiprocessor system, despite the fault containment boundaries between cells, resource sharing is implemented through the cooperation of the various cell kernels, but the policy is implemented outside the kernels in the wax process. Both memory and processors can be shared.

A system called "Cellular IRIX" developed and marketed by Silicon Graphics Inc. Mountain View, Calif., supports modular computing by extending traditional symmetric multiprocessing systems. The Cellular IRIX architecture distributes global kernel text and data into optimized SMP-sized chunks or "cells". Cells represent a control domain consisting of one or more machine modules, where each module consists of processors, memory, and I/O. Applications running on these cells rely extensively on a full set of local operating system services, including local copies of operating system text and kernel data structures, but only one instance of the operating system exists on the entire system.

Inter-cell coordination allows application images to directly and transparently utilize processing, memory and I/O resources from other cells without incurring the overhead of data copies or extra context switches.

Another existing architecture called NUMA-Q developed and marketed by Sequent Computer Systems, Inc., Beaverton, Oreg. uses "quads", or a group of four

processors per portion of memory, as the basic building block for NUMA-Q SMP nodes. Adding I/O to each quad further improves performance. Therefore, the NUMA-Q architecture not only distributes physical memory but puts a predetermined number of processors and PCI slots next to each processor. The memory in each quad is not local memory in the traditional sense. Rather, it is a portion of the physical memory address space and has a specific address range. The address map is divided evenly over memory, with each quad containing a contiguous portion of address space. Only one copy of the operating system is running and, as in any SMP system, it resides in memory and runs processes without distinction and simultaneously on one or more processors. Accordingly, while many attempts have been made at providing a flexible computer system having mechanisms for reconfiguring the system "on the fly" without rebooting, existing systems each have significant shortcomings. Therefore, it would be desirable to have a new computer system design which provides improved flexibility, and resource migration capabilities. Further, it would be desirable to have a new computer system design which permits dynamic reconfiguration under control of the operating systems running on the system and without system administrator intervention.

#### SUMMARY OF THE INVENTION

In accordance with the principles of the present invention, multiple instances of operating systems execute cooperatively in a single multiprocessor computer wherein a single physical machine is subdivided by software into multiple partitions, each with the ability to run a distinct copy, or instance, of an operating system. Each individual instance has the resources it needs to execute independently, but instances cooperate to migrate resources from one partition to another. In accordance with the principles of the invention, the migration can be initiated and carried out under control of the operating system instances "on the fly" without intervention of the system administrator. Alternatively, a system administrator can reconfigure the system.

In accordance with one embodiment, resource migration is carried out under a "push" model in which resources are controlled by an owning partition and must be released by that partition before they can be migrated to another partition. In accordance with this model, a first operating system instance which requires a resource first requests the resource from a second instance. In response to this request, the second instance determines whether it can spare the resource, and if so, begins to bring the resource into an idle state. The resource is transferred when the second instance stops using the resource. Also, in accordance with the "push" model, the request for the resource migration may be initiated from an operator running a program on the second instance or from policy management software which initiates the request.

#### BRIEF DESCRIPTION OF THE DRAWINGS

The above and further advantages of the invention may be better understood by referring to the following description in conjunction with the accompanying drawings and which:

FIG. 1 is a schematic block diagram of a hardware platform illustrating several system building blocks.

FIG. 2 is a schematic diagram of an APMP computer system constructed in accordance with the principles of the present invention illustrating several partitions.

FIG. 3 is a schematic diagram of a configuration tree which represents hardware resource configurations and software configurations and their component parts with child and sibling pointers.

FIG. 4 is a schematic diagram of the configuration tree shown in FIG. 3 and rearranged to illustrate the assignment of hardware to software instances by ownership pointers.

FIG. 5 is a flowchart outlining steps in an illustrative routine for creating an APMP computer system in accordance with the principles of the present invention.

FIG. 6 is a flowchart illustrating the steps in an illustrative routine for creating entries in an APMP system management database which maintains information concerning the APMP system and its configuration.

FIG. 7 forms a flowchart illustrating in detail the steps in an illustrative routine for creating an APMP computer system in accordance with the principles of the present invention.

FIG. 8 forms a flowchart illustrating the steps in an illustrative routine followed by an operating system instance to join an APMP computer system which is already created.

FIG. 9 is a flowchart generally illustrating the steps for migrating a resource from one partition to another.

FIGS. 10A-10E are block schematic illustrations which graphically depict the steps for migrating a resource from one partition to another.

#### DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

A computer platform constructed in accordance with the principles of the present invention is a multi-processor system capable of being partitioned to allow the concurrent execution of multiple instances of operating system software. The system does not require hardware support for the partitioning of its memory, CPUs and I/O subsystems, but some hardware may be used to provide additional hardware assistance for isolating faults, and minimizing the cost of software engineering. The following specification describes the interfaces and data structures required to support the inventive software architecture. The interfaces and data structures described are not meant to imply a specific operating system must be used, or that only a single type of operating system will execute concurrently. Any operating system which implements the software requirements discussed below can participate in the inventive system operation.

#### System Building Blocks

The inventive software architecture operates on a hardware platform which incorporates multiple CPUs, memory and I/O hardware. Preferably, a modular architecture such as that shown in FIG. 1 is used, although those skilled in the art will understand that other architectures can also be used, which architectures need not be modular. FIG. 1 illustrates a computing system constructed of four basic system building blocks (SBBs) 100-106. In the illustrative embodiment, each building block, such as block 100, is identical and comprises several CPUs 108-114, several memory slots (illustrated collectively as memory 120), an I/O processor 118, and a port 116 which contains a switch (not shown) that can connect the system to another such system. However, in other embodiments, the building blocks need not be identical. Large multiprocessor systems can be constructed by connecting the desired number of system building blocks by means of their ports. Switch technology, rather than bus technology, is employed to connect building block components in order to both achieve the improved bandwidth and to allow for non-uniform memory architectures (NUMA).

In accordance with the principles of the invention, the hardware switches are arranged so that each CPU can address all available memory and I/O ports regardless of the number of building blocks configured as schematically illustrated by line 122. In addition, all CPUs may communicate to any or all other CPUs in all SBBs with conventional mechanisms, such as inter-processor interrupts. Consequently, the CPUs and other hardware resources can be associated solely with software. Such a platform architecture is inherently scalable so that large amounts of processing power, memory and I/O will be available in a single computer.

An APMP computer system 200 constructed in accordance with the principles of the present invention from a software view is illustrated in FIG. 2. In this



system, the hardware components have been allocated to allow concurrent execution of multiple operating system instances 208, 210, 212.

In a preferred embodiment, this allocation is performed by a software program called a "console" program, which, as will hereinafter be described in detail, is loaded into memory at power up. Console programs are shown schematically in FIG. 2 as programs 213, 215 and 217. The console program may be a modification of an existing administrative program or a separate program which interacts with an operating system to control the operation of the preferred embodiment. The console program does not virtualize the system resources, that is, it does not create any software layers between the running operating systems 208, 210 and 212 and the physical hardware, such as memory and I/O units (not shown in FIG. 2.) Nor is the state of the running operating systems 208, 210 and 212 swapped to provide access to the same hardware. Instead, the inventive system logically divides the hardware into partitions. It is the responsibility of operating system instance 208, 210, and 212 to use the resources appropriately and provide coordination of resource allocation and sharing. The hardware platform may optionally provide hardware assistance for the division of resources, and may provide fault barriers to minimize the ability of an operating system to corrupt memory, or affect devices controlled by another operating system copy.

The execution environment for a single copy of an operating system, such as copy 208 is called a "partition" 202, and the executing operating system 208 in partition 202 is called "instance" 208. Each operating system instance is capable of booting and running independently of all other operating system instances in the computer system, and can cooperatively take part in sharing resources between operating system instances as described below.

In order to run an operating system instance, a partition must include a hardware restart parameter block (HWRPB), a copy of a console program, some amount of memory, one or more CPUs, and at least one I/O bus which must have a dedicated physical port for the console. The HWRPB is a configuration block which is passed between the console program and the operating system.

Each of console programs 213, 215 and 217, is connected to a console port, shown as ports 214, 216 and 218, respectively. Console ports, such as ports 214, 216 and 218, generally come in the form of a serial line port, or attached graphics, keyboard and mouse options. For the purposes of the inventive computer system, the capability of supporting a dedicated graphics port and associated input devices is not required, although a specific operating system may require it. The base assumption is that a serial port is sufficient for each partition. While a separate terminal, or independent graphics console, could be used to display information generated by each console, preferably the serial lines 220, 222 and 224, can all be connected to a single multiplexer 226 attached to a workstation, PC, or LAT 228 for display of console information.

It is important to note that partitions are not synonymous with system building blocks. For example, partition 202 may comprise the hardware in building blocks 100 and 106 in FIG. 1 whereas partitions 204 and 206 might comprise the hardware in building blocks 102 and 104, respectively. Partitions may also include part of the hardware in a building block.

Partitions can be "initialized" or "uninitialized." An initialized partition has sufficient resources to execute an operating system instance, has a console program image loaded, and a primary CPU available and executing. An initialized partition may be under control of a console program, or may be executing an operating system instance. In an initialized state, a partition has full ownership and control of hardware components assigned to it and only the partition itself may release its components.

In accordance with the principles of the invention, resources can be reassigned from one initialized partition to another. Reassignment of resources can only be performed by the initialized partition to which the resource is currently

assigned. When a partition is in an uninitialized state, other partitions may reassign its hardware components and may delete it.

An uninitialized partition is a partition which has no primary CPU executing either under control of a console program or an operating system. For example, a partition may be uninitialized due to a lack of sufficient resources at power up to run a primary CPU, or when a system administrator is reconfiguring the computer system. When in an uninitialized state, a partition may reassign its hardware components and may be deleted by another partition. Unassigned resources may be assigned by any partition.

Partitions may be organized into "communities" which provide the basis for grouping separate execution contexts to allow cooperative resource sharing. Partitions in the same community can share resources. Partitions that are not within the same community cannot share resources. Resources may only be manually moved between partitions that are not in the same community by the system administrator by de-assigning the resource (and stopping usage), and manually reconfiguring the resource. Communities can be used to create independent operating system domains, or to implement user policy for hardware usage. In FIG. 2, partitions 202 and 204 have been organized into community 230. Partition 206 may be in its own community 205. Communities can be constructed using the configuration tree described below and may be enforced by hardware.

#### The Console Program

When a computer system constructed in accordance with the principles of the present invention is enabled on a platform, multiple HWRPB's must be created, multiple console program copies must be loaded, and system resources must be assigned in such a way that each HWRPB is associated with specific components of the system. To do this, the first console program to run will create a configuration tree structure in memory which represents all of the hardware in the system. The tree will also contain the software partitioning information, and the assignments of hardware to partitions and is discussed in detail below. More specifically, when the APMP system is powered up, a CPU will be selected as a primary CPU in a conventional manner by hardware which is specific to the platform on which the system is running. The primary CPU then loads a copy of a console program into memory. This console copy is called a "master console" program. The primary CPU initially operates under control of the master console program to perform testing and checking assuming that there is a single system which owns the entire machine. Subsequently, a set of environment variables are loaded which define the system partitions. Finally, the master console creates and initializes the partitions based on the environment variables. In this latter process the master console operates to create the configuration tree, to create additional HWRPB data blocks, to load the additional console program copies, and to start the CPUs on the alternate HWRPBs. Each partition then has an operating system instance running on it, which instance cooperates with a console program copy also running in that partition. In an unconfigured APMP system, the master console program will initially create a single partition containing the primary CPU, a minimum amount of memory, and a physical system administrator's console selected in a platform-specific way. Console program commands will then allow the system administrator to create additional partitions, and configure I/O buses, memory, and CPUs for each partition. After associations of resources to partitions have been made by the console program, the associations are stored in non-volatile RAM to allow for an automatic configuration of the system during subsequent boots. During subsequent boots, the master console program must validate the current configuration with the stored configuration to handle the removal and addition of new components. Newly-added components are placed into an unassigned state, until they are assigned by the system administrator. If the removal of a hardware component results in a partition with insufficient resources to run an

operating system, resources will continue to be assigned to the partition, but it will be incapable of running an operating system instance until additional new resources are allocated to it.

As previously mentioned, the console program communicates with an operating system instance by means of an HWRPB which is passed to the operating system during operating system boot up. The fundamental requirements for a console program are that it should be able to create multiple copies of HWRPBs and itself. Each HWRPB copy created by the console program will be capable of booting an independent operating system instance into a private section of memory and each operating system instance booted in this manner can be identified by a unique value placed into the HWRPB. The value indicates the partition, and is also used as the operating system instance ID.

In addition, the console program is configured to provide a mechanism to remove a CPU from the available CPUs within a partition in response to a request by an operating system running in that partition. Each operating system instance must be able to shutdown, halt, or otherwise crash in a manner that control is passed to the console program. Conversely, each operating system instance must be able to reboot into an operational mode, independently of any other operating system instance.

Each HWRPB which is created by a console program will contain a CPU slot-specific database for each CPU that is in the system, or that can be added to the system without powering the entire system down. Each CPU that is physically present will be marked "present", but only CPUs that will initially execute in a specific partition will be marked "available" in the HWRPB for the partition. The operating system instance running on a partition will be capable of recognizing that a CPU may be available at some future time by a present (PP) bit in a per-CPU state flag fields of the HWRPB, and can build data structures to reflect this. When set, the available (PA) bit in the per-CPU state flag fields indicates that the associated CPU is currently associated with the partition, and can be invited to join SMP operation.

#### The Configuration Tree

As previously mentioned, the master console program creates a configuration tree which represents the hardware configuration, and the assignment of each component in the system to each partition. Each console program then identifies the configuration tree to its associated operating system instance by placing a pointer to the tree in the HWRPB.

Referring to FIG. 3, the configuration tree 300 represents the hardware components in the system, the platform constraints and minimums, and the software configuration. The master console program builds the tree using information discovered by probing the hardware, and from information stored in non-volatile RAM which contains configuration information generated during previous initializations.

The master console may generate a single copy of the tree which copy is shared by all operating system instances, or it may replicate the tree for each instance. A single copy of the tree has the disadvantage that it can create a single point of failure in systems with independent memories. However, platforms that generate multiple tree copies require the console programs to be capable of keeping changes to the tree synchronized.

The configuration tree comprises multiple nodes including root nodes, child nodes and sibling nodes. Each node is formed of a fixed header and a variable length extension for overlaid data structures. The tree starts with a tree root node 302 representing the entire system box, followed by branches that describe the hardware configuration (hardware root node 304), the software configuration (software root node 306), and the minimum partition requirements (template root node 308.) In FIG. 3, the arrows represent child and sibling relationships. The children of a node represent component parts of the hardware or software configuration. Siblings represent peers of a component that may not be related

except by having the same parent. Nodes in the tree 300 contain information on the software communities and operating system instances, hardware configuration, configuration constraints, performance boundaries and hot-swap capabilities. The nodes also provide the relationship of hardware to software ownership, or the sharing of a hardware component.

The nodes are stored contiguously in memory and the address offset from the tree root node 302 of the tree 300 to a specific node forms a "handle" which may be used from any operating system instance to unambiguously identify the same component on any operating system instance. In addition, each component in the inventive computer system has a separate ID. This may illustratively be a 64-bit unsigned value. The ID must specify a unique component when combined with the type and subtype values of the component. That is, for a given type of component, the ID must identify a specific component. The ID may be a simple number, for example the CPU ID, it may be some other unique encoding, or a physical address. The component ID and handle allow any member of the computer system to identify a specific piece of hardware or software. That is, any partition using either method of specification must be able to use the same specification, and obtain the same result.

As described above, the inventive computer system is composed of one or more communities which, in turn, are composed of one or more partitions. By dividing the partitions across the independent communities, the inventive computer system can be placed into a configuration in which sharing of devices and memory can be limited. Communities and partitions will have IDs which are densely packed. The hardware platform will determine the maximum number of partitions based on the hardware that is present in the system, as well as having a platform maximum limit. Partition and community IDs will never exceed this value during runtime. IDs will be reused for deleted partitions and communities. The maximum number of communities is the same as the maximum number of partitions. In addition, each operating system instance is identified by a unique instance identifier, for example a combination of the partition ID plus an incarnation number.

The communities and partitions are represented by a software root node 306, which has community node children (of which community node 310 is shown), and partition node grandchildren (of which two nodes, 312 and 314, are shown.) The hardware components are represented by a hardware root node 304 which contains children that represent a hierarchical representation of all of the hardware currently present in the computer system. "Ownership" of a hardware component is represented by a handle in the associated hardware node which points to the appropriate software node (310, 312 or 314.) These handles are illustrated in FIG. 4 which will be discussed in more detail below. Components that are owned by a specific partition will have handles that point to the node representing the partition. Hardware which is shared by multiple partitions (for example, memory) will have handles that point to the community to which sharing is confined. Un-owned hardware will have a handle of zero (representing the tree root node 302).

Hardware components place configuration constraints on how ownership may be divided. A "config" handle in the configuration tree node associated with each component determines if the component is free to be associated anywhere in the computer system by pointing to the hardware root node 304. However, some hardware components may be bound to an ancestor node and must be configured as part of this node. Examples of this are CPUs, which may have no constraints on where they execute, but which are a component part of a system building block (SBB), such as SBBs 322 or 324. In this case, even though the CPU is a child of the SBB, its config handle will point to the hardware root node 304. An I/O bus, however, may not be able to be owned by a partition other than the partition that owns its I/O processor. In this case, the configuration tree node representing the I/O bus would have a config handle pointing to the I/O

processor. Because the rules governing hardware configuration are platform specific, this information is provided to the operating system instances by the config handle.

Each hardware component also has an "affinity" handle. The affinity handle is identical to the config handle, except that it represents a configuration which will obtain the best performance of the component. For example, a CPU or memory may have a config handle which allows it to be configured anywhere in the computer system (it points to the hardware root node 304), however, for optimal performance, the CPU or memory should be configured to use the System Building Block of which they are a part. The result is that the config pointer points to the hardware root node 304, but the affinity pointer points to an SBB node such as node 322 or node 324. The affinity of any component is platform specific, and determined by the firmware. Firmware may use this information when asked to form "optimal" automatic configurations.

Each node also contains several flags which indicate the type and state of the node. These flags include a node\_hotswap flag which indicates that the component represented is a "hot swappable" component and can be powered down independently of its parent and siblings. However, all children of this node must power down if this component powers down. If the children can power down independently of this component, they must also have this bit set in their corresponding nodes. Another flag is a node\_unavailable flag which, when set, indicates that the component represented by the node is not currently available for use. When a component is powered down (or is never powered up) it is flagged as unavailable.

Two flags, node\_hardware and node\_template, indicate the type of node. Further flags, such as node\_initialized and node\_cpu\_primary may also be provided to indicate whether the node represents a partition which has been initialized or a CPU that is currently a primary CPU.

The configuration tree 300 may extend to the level of device controllers, which will allow the operating system to build bus and device configuration tables without probing the buses. However, the tree may also end at any level, if all components below it cannot be configured independently. System software will still be required to probe for bus and device information not provided by the tree.

The console program implements and enforces configuration constraints, if any, on each component of the system. In general, components are either assignable without constraints (for example, CPUs may have no constraints), or are configurable only as a part of another component (a device adapter, for example, may be configurable only as a part of its bus). A partition which is, as explained above, a grouping of CPUs, memory, and I/O devices into a unique software entity also has minimum requirements. For example, the minimum hardware requirements for a partition are at least one CPU, some private memory (platform dependent minimum, including console memory) and an I/O bus, including a physical, non-shared, console port.

The minimal component requirements for a partition are provided by the information contained in the template root node 308. The template root node 308 contains nodes, 316, 318 and 320, representing the hardware components that must be provided to create a partition capable of execution of a console program and an operating system instance. Configuration editors can use this information as the basis to determine what types, and how many resources must be available to form a new partition.

During the construction of a new partition, the template subtree will be "walked", and, for each node in the template subtree, there must be a node with the same type and subtype owned by the new partition so that it will be capable of loading a console program and booting an operating system instance. If there are more than one node of the same type and subtype in the template tree, there must also be multiple nodes in the new partition. The console program will use

the template to validate that a new partition has the minimum requirements prior to attempting to load a console program and initialize operation. The following is a detailed example of a particular implementation of configuration tree nodes. It is intended for descriptive purposes only and is not intended to be limiting. Each HWRPB must point to a configuration tree which provides the current configuration, and the assignments of components to partitions. A configuration pointer (in the CONFIG field) in the HWRPB is used to point to the configuration tree. The CONFIG field points to a 64-byte header containing the size of the memory pool for the tree, and the initial checksum of the memory. Immediately following the header is the root node of the tree. The header and root node of the tree will be page aligned. The total size in bytes of the memory allocated for the configuration tree is located in the first quadword of the header. The size is guaranteed to be in multiples of the hardware page size. The second quadword of the header is reserved for a checksum. In order to examine the configuration tree, an operating system instance maps the tree into its local address space. Because an operating system instance may map this memory with read access allowed for all applications, some provision must be made to prevent a non-privileged application from gaining access to console data to which it should not have access. Access may be restricted by appropriately allocating memory. For example, the memory may be page aligned and allocated in whole pages. Normally, an operating system instance will map the first page of the configuration tree, obtain the tree size, and then remap the memory allocated for configuration tree usage. The total size may include additional memory used by the console for dynamic changes to the tree. Preferably, configuration tree nodes are formed with fixed headers, and may optionally contain type-specific information following the fixed portion. The size field contains the full length of the node, nodes are illustratively allocated in multiples of 64-bytes and padded as needed. The following description defines illustrative fields in the fixed header for a node:

```
typedef struct_gct_node {
    unsigned char    type;
    unsigned char    subtype;
    uint16           size;
    GCT_HANDLE       owner;
    GCT_HANDLE       current_owner;
    GCT_ID           id;
    union {
        uint64       node_flags;
        struct {
            unsigned  node_hardware      : 1;
            unsigned  node_hotswap       : 1;
            unsigned  node_unavailable   : 1;
            unsigned  node_hw_template  : 1;
            unsigned  node_initialized   : 1;
            unsigned  node_cpu_primary   : 1;
        } flag_bits
    } flag_union;
    GCT_HANDLE       config;
    GCT_HANDLE       affinity;
    GCT_HANDLE       parent;
};
```

```

GCT_HANDLE    next_sib;
GCT_HANDLE    prev_sib;
GCT_HANDLE    child;
GCT_HANDLE    reserved;
uint32        magic
} GCT_NODE;

```

In the above definition the type definitions "uint" are unsigned integers with the appropriate bit lengths. As previously mentioned, nodes are located and identified by a handle (identified by the typedef GCT\_HANDLE in the definition above). An illustrative handle is a signed 32-bit offset from the base of the configuration tree to the node. The value is unique across all partitions in the computer system. That is, a handle obtained on one partition must be valid to lookup a node, or as an input to a console callback, on all partitions. The magic field contains a predetermined bit pattern which indicates that the node is actually a valid node.

The tree root node represents the entire system. Its handle is always zero. That is, it is always located at the first physical location in the memory allocated for the configuration tree following the config header. It has the following definition:

```

typedef struct_gct_root_node {
    GCT_NODE    hd;
    uint64      lock;
    uint64      transient_level;
    uint64      current_level;
    uint64      console_req;
    uint64      min_alloc;
    uint64      min_align;
    uint64      base_alloc;
    uint64      base_align;
    uint64      max_phys_address;
    uint64      mem_size;
    uint64      platform_type;
    int32       platform_name;
    GCT_HANDLE  primary_instance;
    GCT_HANDLE  first_free;
    GCT_HANDLE  high_limit;
    GCT_HANDLE  lookaside;
    GCT_HANDLE  available;
    uint32      max_partition;
    int32       partitions;
    int32       communities;
    uint32      max_platform_partition;
    uint32      max_fragments;
    uint32      max_desc;
    char        APMP_id[16];
    char        APMP_id_pad[4];
    int32       bindings;
} GCT_ROOT_NODE;

```

The fields in the root node are defined as follows:

lock

This field is used as a simple lock by software wishing to inhibit changes to the structure of the tree, and the software configuration. When this value is -1 (all bits on) the tree is unlocked; when the value is  $\geq 0$  the tree is locked. This field is modified using atomic operations. The caller of the lock routine passes a partition ID which is written to the lock field. This can be used to assist in fault tracing, and recovery during crashes.

`transient_level`

This field is incremented at the start of a tree update.

`current_level`

This field is updated at the completion of a tree update.

`console_req`

This field specifies the memory required in bytes for the console in the base memory segment of a partition.

`min_alloc`

This field holds the minimum size of a memory fragment, and the allocation unit (fragments size must be a multiple of the allocation). It must be a power of 2.

`min_align`

This field holds the alignment requirements for a memory fragment. It must be a power of 2.

`base_alloc`

This field specifies the minimum memory in bytes (including `console_req`) needed for the base memory segment for a partition. This is where the console, console structures, and operating system will be loaded for a partition. It must be greater or equal to `minAlloc` and a multiple of `minAlloc`.

`base_align`

This field holds the alignment requirement for the base memory segment of a partition. It must be a power of 2, and have an alignment of at least

`min_align`.

`max_phys_address`

The field holds the calculated largest physical address that could exist on the system, including memory subsystems that are not currently powered on and available.

`mem_size`

This field holds the total memory currently in system.

`platform_type`

This field stores the type of platform taken from a field in the HWRPBB.

`platform_name`

This field holds an integer offset from the base of the tree root node to a string representing the name of the platform.

`primary_instance`

This field stores the partition ID of the first operating system instance.

`first_free`

This field holds the offset from the tree root node to the first free byte of memory pool used for new nodes.

`high_limit`

This field holds the highest address at which a valid node can be located within the configuration tree. It is used by callbacks to validate that a handle is legal.

`lookaside`

This field is the handle of a linked list of nodes that have been deleted, and that may be reclaimed. When a community or partition are deleted, the node is linked into this list, and creation of a new partition or community will look at this list before allocating from free pool.

`available`

This field holds the number of bytes remaining in the free pool pointed to by the `first_free` field.

`max_partitions`

This field holds the maximum number of partitions computed by the platform based on the amount of hardware resources currently available.

`partitions`

This field holds an offset from the base of the root node to an array of handles.



Each partition ID is used as an index into this array, and the partition node handle is stored at the indexed location. When a new partition is created, this array is examined to find the first partition ID which does not have a corresponding partition node handle and this partition ID is used as the ID for the new partition.

communities

This field also holds an offset from the base of the root node to an array of handles. Each community ID is used as an index into this array, and a community node handle is stored in the array. When a new community is created, this array is examined to find the first community ID which does not have a corresponding community node handle and this community ID is used as the ID for the new community. There cannot be more communities than partitions, so the array is sized based on the maximum number of partitions.

max\_platform\_partition

This field holds the maximum number of partitions that can simultaneously exist on the platform, even if additional hardware is added (potentially inswapped).

max\_fragments

This field holds a platform defined maximum number of fragments into which a memory descriptor can be divided. It is used to size the array of fragments in the memory descriptor node.

max\_desc

This field holds the maximum number of memory descriptors for the platform.

APMP\_id

This field holds a system ID set by system software and saved in non-volatile RAM.

APMP\_id\_pad

This field holds padding bytes for the APMP ID.

bindings

This field holds an offset to an array of "bindings" Each binding entry describes a type of hardware node, the type of node the parent must be, the configuration binding, and the affinity binding for a node type. Bindings are used by software to determine how node types are related and configuration and affinity rules.

A community provides the basis for the sharing of resources between partitions. While a hardware component may be assigned to any partition in a community, the actual sharing of a device, such as memory, occurs only within a community. The community node 310 contains a pointer to a control section, called an APMP database, which allows the operating system instances to control access and membership in the community for the purpose of sharing memory and communications between instances. The APMP database and the creation of communities are discussed in detail below. The configuration ID for the community is a signed 16-bit integer value assigned by the console program. The ID value will never be greater than the maximum number of partitions that can be created on the platform.

A partition node, such as node 312 or 314, represents a collection of hardware that is capable of running an independent copy of the console program, and an independent copy of an operating system. The configuration ID for this node is a signed 16-bit integer value assigned by the console. The ID will never be greater than the maximum number of partitions that can be created on the platform. The node has the definition:

```
typedef struct gct_partition_node {
    GCT_NODE    hd;
    uint64      hwrpb;
    uint64      incarnation;
    uint64      priority;
    int32       os_type;
```

```

        uint32      partition_reserved_1;
        uint64      instance_name_format;
        char        instance_name[128];
    } GCT_PARTITION_NODE;

```

The defined fields have the definitions:

**hwrpb**

This field holds the physical address of the hardware restart parameter block for this partition. To minimize changes to the HWRPB, the HWRPB does not contain a pointer to the partition, or the partition ID. Instead, the partition nodes contain a pointer to the HWRPB. System software can then determine the partition ID of the partition in which it is running by searching the partition nodes for the partition which contains the physical address of its HWRPB.

**incarnation**

This field holds a value which is incremented each time the primary CPU of the partition executes a boot or restart operation on the partition.

**priority**

This field holds a partition priority.

**os\_type**

This field holds a value which indicates the type of operating system that will be loaded in the partition.

**partition\_reserved\_1**

This field is reserved for future use.

**instance\_name\_format**

This field holds a value that describes the format of the instance name string.

**instance\_name**

This field holds a formatted string which is interpreted using the `instance_name_format` field. The value in this field provides a high-level path name to the operating system instance executing in the partition. This field is loaded by system software and is not saved across power cycles. The field is cleared at power up and at partition creation and deletion.

A System Building Block node, such as node 322 or 324, represents an arbitrary piece of hardware, or conceptual grouping used by system platforms with modular designs such as that illustrated in FIG. 2. A QBB (Quad Building Block) is a specific example of an SBB and corresponds to units such as units 100, 102, 104 and 106 in FIG. 1. Children of the SBB nodes 322 and 324 include input/output processor nodes 326 and 340.

CPU nodes, such as nodes 328-332 and 342-346, are assumed to be capable of operation as a primary CPU for SMP operation. In the rare case where a CPU is not primary capable, it will have a SUBTYPE code indicating that it cannot be used as a primary CPU in SMP operation. This information is critical when configuring resources to create a new partition. The CPU node will also carry information on where the CPU is currently executing. The primary for a partition will have the `NODE_CPU_PRIMARY` flag set in the `NODE_FLAGS` field. The CPU node has the following definition:

```

typedef struct gct_cpu_node {
    GCT_NODE    hd;
} GCT_CPU_NODE;

```

A memory subsystem node, such as node 334 or 348, is a "pseudo" node that groups together nodes representing the physical memory controllers and the assignments of the memory that the controllers provide. The children of this node consist of one or more memory controller nodes (such as nodes 336 and 350) which the console has configured to operate together (interleaved), and one or more memory descriptor nodes (such as nodes 338 and 352) which describe physically contiguous ranges of memory.

A memory controller node (such as nodes 336 or 350) is used to express a physical hardware component, and its owner is typically the partition which will handle errors, and initialization. Memory controllers cannot be assigned

to communities, as they require a specific operating system instance for initialization, testing and errors. However, a memory description, defined by a memory descriptor node, may be split into "fragments" to allow different partitions or communities to own specific memory ranges within the memory descriptor. Memory is unlike other hardware resources in that it may be shared concurrently, or broken into "private" areas. Each memory descriptor node contains a list of subset ranges that allow the memory to be divided among partitions, as well as shared between partitions (owned by a community). A memory descriptor node (such as nodes 338 or 352) is defined as:

```
typedef struct gct_mem_desc_node {
    GCT_NODE      hd;
    GCT_MEM_INFO  mem_info;
    int32         mem_frag;
} GCT_MEM_DESC_NODE;
```

The mem\_info structure has the following definition:

```
typedef struct gct_mem_info {
    uint64      base_pa;
    uint64      base_size;
    uint32      desc_count;
    uint32      info_fill;
} GCT_MEM_INFO;
```

The mem\_frag field holds an offset from the base of the memory descriptor node to an array of GCT\_MEM\_DESC structures which have the definition:.

```
typedef struct gct_mem_desc {
    uint64      pa;
    uint64      size;
    GCT_HANDLE  mem_owner;
    GCT_HANDLE  mem_current_owner;
    union {
        uint32      mem_flags;
        struct {
            unsigned mem_console   : 1;
            unsigned mem_private   : 1;
            unsigned mem_shared    : 1;
            unsigned base          : 1;
#define CGT_MEM_CONSOLE 0.times.1
#define CGT_MEM_PRIVATE 0.times.2
#define CGT_MEM_SHARED  0.times.4
#define CGT_MEM_CONSOLE 0.times.8
        } flag_bits;
    } flag_union;
    uint32      mem_fill;
} GCT_MEM_DESC;
```

The number of fragments in a memory description node (nodes 338 or 352) is limited by platform firmware. This creates an upper bound on memory division, and limits unbounded growth of the configuration tree. Software can determine the maximum number of fragments from the max\_fragments field in the tree root node 302 (discussed above), or by calling an appropriate console callback function to return the value. Each fragment can be assigned to any partition, provided that the config binding, and the ownership of the memory descriptor and memory subsystem nodes allow it. Each fragment contains a base physical address, size, and owner field, as well as flags indicating the type of usage. To allow shared memory access, the memory subsystem parent node, and the memory descriptor node must be owned by a community. The fragments within the memory descriptor may then be owned by the community (shared) or by any partition within the community.

Fragments can have minimum allocation sizes and alignments provided in the tree root node 302. The base memory for a partition (the fragments where the console and operating system will be loaded) may have a greater allocation and alignment than other fragments (see the tree root node definition above). If the owner field of the memory descriptor node is a partition, then the fragments can only be owned by that partition.

FIG. 4 illustrates the configuration tree shown in FIG. 3 when it is viewed from a perspective of ownership. The console program for a partition relinquishes ownership and control of the partition resources to the operating system instance running in that partition when the primary CPU for that partition starts execution. The concept of "ownership" determines how the hardware resources and CPUs are assigned to software partitions and communities. The configuration tree has ownership pointers illustrated in FIG. 4 which determine the mapping of hardware devices to software such as partitions (exclusive access) and communities (shared access). An operating system instance uses the information in the configuration tree to determine to which hardware resources it has access and reconfiguration control.

Passive hardware resources which have no owner are unavailable for use until ownership is established. Once ownership is established by altering the configuration tree, the operating system instances may begin using the resources. When an instance makes an initial request, ownership can be changed by causing the owning operating system to stop using a resource or by a console program taking action to stop using a resource in a partition where no operating system instance is executing. The configuration tree is then altered to transfer ownership of the resource to another operating system instance. The action required to cause an operating system to stop using a hardware resource is operating system specific, and may require a reboot of the operating system instances affected by the change.

To manage the transition of a resource from an owned and active state, to a unowned and inactive state, two fields are provided in each node of the tree. The owner field represents the owner of a resource and is loaded with the handle of the owning software partition or community. At power up of an APMP system, the owner fields of the hardware nodes are loaded from the contents of non-volatile RAM to establish an initial configuration.

To change the owner of a resource, the handle value is modified in the owner field of the hardware component, and in the owner fields of any descendants of the hardware component which are bound to the component by their config handles. The current\_owner field represents the current user of the resource. When the owner and current\_owner fields hold the same non-zero value, the resource is owned and active. Only the owner of a resource can de-assign the resource (set the owner field to zero). A resource that has null owner and current\_owner fields is unowned, and inactive. Only resources which have null owner and current\_owner fields may be assigned to a new partition or community. When a resource is de-assigned, the owner may decide to de-assign the owner field, or both the owner and current\_owner fields. The decision is based on the ability of the owning operating system instance running in the partition to discontinue the use of the resource prior to de-assigning ownership. In the case where a reboot is required to relinquish ownership, the owner field is cleared, but the current\_owner field is not changed. When the owning operating system instance reboots, the console program can clear any current\_owner fields for resources that have no owner during initialization.

During initialization, the console program will modify the current\_owner field to match the owner field for any node of which it is the owner, and for which the current\_owner field is null. System software should only use hardware of which it is the current owner. In the case of a de-assignment of a resource which is owned by a community, it is the responsibility of system software to manage the transition between states. In some embodiments, a resource may be

loaned to another partition. In this condition, the owner and current\_owner fields are both valid, but not equal. The following table summarizes the possible resource states and the values of the owner and current\_owner fields:

TABLE 1

| owner field | value                 | current_owner field | value | Resource State            |
|-------------|-----------------------|---------------------|-------|---------------------------|
| none        | none                  |                     |       | unowned, and inactive     |
| none        | valid                 |                     |       | unowned, but still active |
| valid       | none                  |                     |       | owned, not yet active     |
| valid       | equal to owner        |                     |       | owned and active          |
| valid       | is not equal to owner |                     |       | loaned                    |

Because CPUs are active devices, and sharing of CPUs means that a CPU could be executing in the context of a partition which may not be its "owner", ownership of a CPU is different from ownership of a passive resource. The CPU node in the configuration tree provides two fields that indicate which partition a CPU is nominally "owned" by, and in which partition the CPU is currently executing. The owner field contains a value which indicates the nominal ownership of the CPU, or more specifically, the partition in which the CPU will initially execute at system power up.

Until an initial ownership is established (that is, if the owner field is unassigned), CPUs are placed into a HWRPB context decided by the master console, but the HWRPB available bit for the CPU will not be set in any HWRPB. This combination prevents the CPU from joining any operating system instance in SMP operation. When ownership of a CPU is established (the owner field is filled in with a valid partition handle), the CPU will migrate, if necessary, to the owning partition, set the available bit in the HWRPB associated with that partition, and request to join SMP operation of the instance running in that partition, or join the console program in SMP mode. The combination of the present and available bits in the HWRPB tell the operating system instance that the CPU is available for use in SMP operation, and the operating system instance may use these bits to build appropriate per-CPU data structures, and to send a message to the CPU to request it to join SMP operation.

When a CPU sets the available bit in an HWRPB, it also enters a value into the current\_owner field in its corresponding CPU node in the configuration tree. The current\_owner field value is the handle of the partition in which the CPU has set the active HWRPB bit and is capable of joining SMP operation. The current\_owner field for a CPU is only set by the console program. When a CPU migrates from one partition to another partition, or is halted into an unassigned state, the current\_owner field is cleared (or changed to the new partition handle value) at the same time that the available bit is cleared in the HWRPB. The current\_owner field should not be written to directly by system software, and only reflects which HWRPB has the available bit set for the CPU. **During runtime, an operating system instance can temporarily "loan" a CPU to another partition without changing the nominal ownership of the CPU. The traditional SMP concept of ownership using the HWRPB present and available bits is used to reflect the current execution context of the CPU by modifying the HWRPB and the configuration tree in atomic operations. The current\_owner field can further be used by system software in one of the partitions to determine in which partition the CPU is currently executing (other instances can determine the location of a particular CPU by examining the configuration tree.)**

It is also possible to de-assign a CPU and return it into a state in which the available bit is not set in any HWRPB, and the current\_owner field in the configuration tree node for the CPU is cleared. This is accomplished by halting the execution of the CPU and causing the console program to clear the owner field in the configuration tree node, as well as the current\_owner field and the available HWRPB bit. The CPU will then execute in console mode and poll the owner field waiting for a valid partition handle to be written to it. System

software can then establish a new owner, and the CPU begin execution in the new partition.

Illustrative ownership pointers are illustrated in FIG. 4 by arrows. Each of the nodes in FIG. 4 that corresponds to a similar node in FIG. 3 is given a corresponding number. For example, the software root node denoted in FIG. 3 as node 306 is denoted as node 406 in FIG. 4. As shown in FIG. 4, the community 410 is "owned" by the software root 406. Likewise, the system building blocks 1 and 2 (422 and 425) are owned by the community 410. Similarly, partitions 412 and 414 are also owned by the community 410.

Partition 412 owns CPUs 428-432 and the I/O processor 426. The memory controller 436 is also a part of partition 1 (412). In a like manner, partition 2 (414) owns CPUs 442-446, I/O processor 440 and memory controller 450.

The common or shared memory in the system is comprised of memory subsystems 434 and 448 and memory descriptors 438 and 452. These are owned by the community 410. Thus, FIG. 4 describes the layout of the system as it would appear to the operating system instances.

#### Operating System Characteristics

As previously mentioned, the illustrative computer system can operate with several different operating systems in different partitions. However, conventional operating systems may need to be modified in some aspects in order to make them compatible with the inventive system, depending on how the system is configured. Some sample modifications for the illustrative embodiment are listed below:

1. Instances may need to be modified to include a mechanism for choosing a "primary" CPU in the partition to run the console and be a target for communication from other instances. The selection of a primary CPU can be done in a conventional manner using arbitration mechanisms or other conventional devices.
2. Each instance may need modifications that allow it to communicate and cooperate with the console program which is responsible for creating a configuration data block that describes the resources available to the partition in which the instance is running. For example, the instance should not probe the underlying hardware to determine what resources are available for usage by the instance. Instead, if it is passed a configuration data block that describes what resources that instance is allowed to access, it will need to work with the specified resources.
3. An instance may need to be capable of starting at an arbitrary physical address and may not be able to reserve any specific physical address in order to avoid conflicting with other operating systems running at that particular address.
4. An instance may need to be capable of supporting multiple arbitrary physical holes in its address space, if it is part of a system configuration in which memory is shared between partitions. In addition, an instance may need to deal with physical holes in its address space in order to support "hot inswap" of memory.
5. An instance may need to pass messages and receive notifications that new resources are available to partitions and instances. More particularly, a protocol is needed to inform an instance to search for a new resource. Otherwise, the instance may never realize that the resource has arrived and is ready for use.
6. An instance may need to be capable of running entirely within its "private memory" if it is used in a system where instances do not share memory. Alternatively, an instance may need to be capable of using physical "shared memory" for communicating or sharing data with other instances running within the computer if the instance is part of a system in which memory is shared. In such a shared memory system, an instance may need to be capable of mapping physical "shared memory" as identified in the configuration tree into its

virtual address space, and the virtual address spaces of the "processes" running within that operating system instance.

7. Each instance may need some mechanism to contact another CPU in the computer system in order to communicate with it.

8. An instance may also need to be able to recognize other CPUs that are compatible with its operations, even if the CPUs are not currently assigned to its partition. For example, the instance may need to be able to ascertain CPU parameters, such as console revision number and clock speed, to determine whether it could run with that CPU, if the CPU was re-assigned to the partition in which the instance is running.

#### Changing the Configuration Tree

**Each console program provides a number of callback functions to allow the associated operating system instance to change the configuration of the APMP system, for example, by creating a new community or partition, or altering the ownership of memory fragments. In addition, other callback functions provide the ability to remove a community, or partition, or to start operation on a newly-created partition.**

However, callback functions do not cause any changes to take place on the running operating system instances. Any changes made to the configuration tree must be acted upon by each instance affected by the change. The type of action that must take place in an instance when the configuration tree is altered is a function of the type of change, and the operating system instance capabilities. For example, moving an input/output processor from one partition to another may require both partitions to reboot. Changing the memory allocation of fragments, on the other hand, might be handled by an operating system instance without the need for a reboot.

**Configuration of an APMP system entails the creation of communities and partitions, and the assignment of unassigned components. When a component is moved from one partition to another, the current owner removes itself as owner of the resource and then indicates the new owner of the resource. The new owner can then use the resource. When an instance running in a partition releases a component, the instance must no longer access the component. This simple procedure eliminates the complex synchronization needed to allow blind stealing of a component from an instance, and possible race conditions in booting an instance during a reconfiguration.**

Once initialized, configuration tree nodes will never be deleted or moved, that is, their handles will always be valid. Thus, hardware node addresses may be cached by software. Callback functions which purport to delete a partition or a community do not actually delete the associated node, or remove it from the tree, but instead flag the node as UNAVAILABLE, and clear the ownership fields of any hardware resource that was owned by the software component.

In order to synchronize changes to the configuration tree, the root node of the tree maintains two counters (`transient_level` and `current_level`). The `transient_level` counter is incremented at the start of an update to the tree, and the `current_level` counter is incremented when the update is complete.

Software may use these counters to determine when a change has occurred, or is occurring to the tree. When an update is completed by a console, an interrupt can be generated to all CPUs in the APMP system. This interrupt can be used to cause system software to update its state based on changes to the tree.

#### Creation of an APMP Computer System

FIG. 5 is a flowchart that illustrates an overview of the formation of the illustrative adaptively-partitioned, multi-processor (APMP) computer system. The routine starts in step 500 and proceeds to step 502 where a master console program is started. If the APMP computer system is being created on power up, the CPU on which the master console runs is chosen by a predetermined mechanism, such as arbitration, or another hardware mechanism. If the APMP computer system is being created on hardware that is already running, a CPU in

the first partition that tries to join the (non-existent) system runs the master console program, as discussed below.

Next, in step 504, the master console program probes the hardware and creates the configuration tree in step 506 as discussed above. If there is more than one partition in the APMP system on power up, each partition is initialized and its console program is started (step 508).

Finally, an operating system instance is booted in at least one of the partitions as indicated in step 510. The first operating system instance to boot creates an APMP database and fills in the entries as described below. APMP databases store information relating to the state of active operating system instances in the system. The routine then finishes in step 512. It should be noted that an instance is not required to participate in an APMP system. The instance can choose not to participate or to participate at a time that occurs well after boot. Those instances which do participate form a "sharing set." The first instance which decides to join a sharing set must create it. There can be multiple sharing sets operating on a single APMP system and each sharing set has its own APMP database.

#### Deciding to Create a New APMP System or to Join an Existing APMP System

An operating system instance running on a platform which is also running the APMP computer system does not necessarily have to be a member of the APMP computer system. The instance can attempt to become a member of the APMP system at any time after booting. This may occur either automatically at boot, or after an operator-command explicitly initiates joining. After the operating system is loaded at boot time, the operating system initialization routine is invoked and examines a stored parameter to see whether it specifies immediate joining and, if so, the system executes a joining routine which is part of the APMP computer system. An operator command would result in an execution of the same routine.

#### APMP Database

An important data structure supporting the inventive software allocation of resources is the APMP database which keeps track of operating system instances which are members of a sharing set. The first operating system instance attempting to set up the APMP computer system initializes an APMP database, thus creating, or instantiating, the inventive software resource allocations for the initial sharing set. Later instances wishing to become part of the sharing set join by registering in the APMP database associated with that sharing set. The APMP database is a shared data structure containing the centralized information required for the management of shared resources of the sharing set. An APMP database is also initialized when the APMP computer system is re-formed in response to an unrecoverable error.

More specifically, each APMP database is a three-part structure. The first part is a fixed-size header portion including basic synchronization structures for creation of the APMP computer system, address-mapping information for the database and offsets to the service-specific segments that make up the second portion. The second portion is an array of data blocks with one block assigned to each potential instance. The data blocks are called "node blocks." The third portion is divided into segments used by each of the computer system sub-facilities. Each sub-facility is responsible for the content of, and synchronizing access to, its own segment.

The initial, header portion of an APMP database is the first part of the APMP database mapped by a joining operating system instance. Portions of the header are accessed before the instance has joined the sharing set, and, in fact, before the instance knows that the APMP computer system exists.

The header section contains:

1. a membership and creation synchronization quadword
2. a computer system software version
3. state information, creation time, incarnation count, etc.



4. a pointer (offset) to a membership mask
5. crashing instance, crash acknowledge bits, etc.
6. validation masks, including a bit for each service
7. memory mapping information (page frame number information) for the entire APMP database
8. offset/length pairs describing each of the service segments (lengths in bytes rounded to pages and offsets full pages) including:
  - shared memory services
  - cpu communications services
  - membership services (if required)
  - locking services

The array of node blocks is indexed by a system partition id (one per instance possible on the current platform) and each block contains:

- instance software version
- interrupt reason mask
- instance state
- instance incarnation
- instance heartbeat
- instance membership timestamp
- little brother instance id and inactive-time; big brother instance id
- instance validation done bit.

An APMP database is stored in shared memory. The initial fixed portion of N physically contiguous pages occupies the first N pages of one of two memory ranges allocated by the first instance to join during initial partitioning of the hardware. The instance directs the console to store the starting physical addresses of these ranges in the configuration tree. The purpose of allocating two ranges is to permit failover in case of hardware memory failure. Memory management is responsible for mapping the physical memory into virtual address space for the APMP database.

The detailed actions taken by an operating system instance are illustrated in FIG. 6. More specifically, when an operating system instance wishes to become a member of a sharing set, it must be prepared to create the APMP computer system if it is the first instance attempting to "join" a non-existent system. In order for the instance to determine whether an APMP system already exists, the instance must be able to examine the state of shared memory as described above. Further, it must be able to synchronize with other instances which may be attempting to join the APMP system and the sharing set at the same time to prevent conflicting creation attempts. The master console creates the configuration tree as discussed above. Subsequently, a region of memory is initialized by the first, or primary, operating system instance to boot, and this memory region can be used for an APMP database.

#### Mapping the APMP Database Header

The goal of the initial actions taken by all operating system instances is to map the header portion of the APMP database and initialize primitive inter-instance interrupt handling to lay the groundwork for a create or join decision. The routine used is illustrated in FIG. 6 which begins in step 600. The first action taken by each instance (step 602) is to engage memory management to map the initial segment of the APMP database as described above. At this time, the array of node blocks in the second database section is also mapped. Memory management maps the initial and second segments of the APMP database into the primary operating system address space and returns the start address and length. The instance then informs the console to store the location and size of the segments in the configuration tree.

Next, in step 604, the initial virtual address of the APMP database is used to allow the initialization routine to zero interrupt reason masks in the node block assigned to the current instance.

A zero initial value is then stored to the heartbeat field for the instance in the node block, and other node block fields. In some cases, the instance attempting to create a new APMP computer system was previously a member of an APMP system and did not withdraw from the APMP system. If this instance is rebooting before the other instances have removed it, then its bit will still be "on" in the system membership mask. Other unusual or error cases can also lead to "garbage" being stored in the system membership mask.

Next, in step 608, the virtual address (VA) of the APMP database is stored in a private cell which is examined by an inter-processor interrupt handler. The handler examines this cell to determine whether to test the per-instance interrupt reason mask in the APMP database header for work to do. If this cell is zero, the APMP database is not mapped and nothing further is done by the handler. As previously discussed, the entire APMP database, including this mask, is initialized so that the handler does nothing before the address is stored. In addition, a clock interrupt handler can examine the same private cell to determine whether to increment the instance-specific heartbeat field for this instance in the appropriate node block. If the private cell is zero, the interrupt handler does not increment the heartbeat field.

At this point, the routine is finished (step 610) and the APMP database header is accessible and the joining instance is able to examine the header and decide whether the APMP computer system does not exist and, therefore, the instance must create it, or whether the instance will be joining an already-existing APMP system.

Once the APMP header is mapped, the header is examined to determine whether an APMP computer system is up and functioning, and, if not, whether the current instance should initialize the APMP database and create the APMP computer system. The problem of joining an existing APMP system becomes more difficult, for example, if the APMP computer system was created at one time, but now has no members, or if the APMP system is being reformed after an error. In this case, the state of the APMP database memory is not known in advance, and a simple memory test is not sufficient. An instance that is attempting to join a possibly existing APMP system must be able to determine whether an APMP system exists or not and, if it does not, the instance must be able to create a new APMP system without interference from other instances. This interference could arise from threads running either on the same instance or on another instance. In order to prevent such interference, the create/join decision is made by first locking the APMP database and then examining the APMP header to determine whether there is a functioning APMP computer system. If there is a properly functioning APMP system, then the instance joins the system and releases the lock on the APMP database. Alternatively, if there is no APMP system, or if there is an APMP system, but it is non-functioning, then the instance creates a new APMP system, with itself as a member and releases the lock on the APMP database.

If there appears to be an APMP system in transition, then the instance waits until the APMP system is again operational or dead, and then proceeds as above. If a system cannot be created, then joining fails.

#### Creating a new APMP Computer System

Assuming that a new APMP system must be created, the creator instance is responsible for allocating the rest of the APMP database, initializing the header and invoking system services. Assuming the APMP database is locked as described above, the following steps are taken by the creator instance to initialize the APMP system (these steps are shown in FIG. 7):

Step 702 the creator instance sets the APMP system state and its node block state to "initializing."

Step 704 the creator instance calls a size routine for each system service with the address of its length field in the header.

Step 706 the resulting length fields are summed and the creator instance calls memory management to allocate space for the entire APMP database by creating a new mapping and deleting the old mapping.

Step 708 the creator instance fills in the offsets to the beginnings of each system service segment.

Step 710 the initialization routine for each service is called with the virtual addresses of the APMP database, the service segment and the segment length.

Step 712 the creator instance initializes a membership mask to make itself the sole member and increments an incarnation count. It then sets creation time, software version, and other creation parameters.

Step 714 the instance then sets itself as its own big and little brother (for heartbeat monitoring purposes as described below).

Step 716 the instance then fills in its instance state as "member" and the APMP system state as "operational."

Step 718 finally, the instance releases the APMP database lock.

The routine then ends in step 720.

#### Joining an Existing APMP Computer System

Assuming an instance has the APMP database locked, the following steps are taken by the instance to become a member of an existing APMP system (shown in FIG. 8):

Step 802 the instance checks to make sure that its instance name is unique. If another current member has the instance's proposed name, joining is aborted.

Step 804 the instance sets the APMP system state and its node block state to "instance joining"

Step 806 the instance calls a memory management routine to map the variable portion of the APMP database into its local address space.

Step 808 the instance calls system joining routines for each system service with the virtual addresses of the APMP database and its segment and its segment length.

Step 810 if all system service joining routines report success, then the instance joining routine continues. If any system service join routine fails, the instance joining process must start over and possibly create a new APMP computer system.

Step 812 assuming that success was achieved in step 810, the instance adds itself to the system membership mask.

Step 814 the instance selects a big brother to monitor its instance health as set forth below.

Step 816 the instance fills in its instance state as "member" and sets a local membership flag.

Step 818 the instance releases the configuration database lock.

The routine then ends in step 820.

The loss of an instance, either through inactivity timeout or a crash, is detected by means of a "heartbeat" mechanism implemented in the APMP database. Instances will attempt to do minimal checking and cleanup and notify the rest of the APMP system during an instance crash. When this is not possible, system services will detect the disappearance of an instance via a software heartbeat mechanism. In particular, a "heartbeat" field is allocated in the APMP database for each active instance. This field is written to by the corresponding instance at time intervals that are less than a predetermined value, for example, every two milliseconds.

Any instance may examine the heartbeat field of any other instance to make a direct determination for some specific purpose. An instance reads the heartbeat field of another instance by reading its heartbeat field twice separated by a two millisecond time duration. If the heartbeat is not incremented between the two reads, the instance is considered inactive (gone, halted at control-P, or hung at or above clock interrupt priority level.) If the instance remains

inactive for a predetermined time, then the instance is considered dead or disinterested.

In addition, a special arrangement is used to monitor all instances because it is not feasible for every instance to watch every other instance, especially as the APMP system becomes large. This arrangement uses a "big brother--little brother" scheme. More particularly, when an instance joins the APMP system, before releasing the lock on the APMP database, it picks one of the current members to be its big brother and watch over the joining instance. The joining instance first assumes big brother duties for its chosen big brother's current little brother, and then assigns itself as the new little brother of the chosen instance. Conversely, when an instance exits the APMP computer system while still in operation so that it is able to perform exit processing, and while it is holding the lock on the APMP database, it assigns its big brother duties to its current big brother before it stops incrementing its heartbeat.

Every clock tick, after incrementing its own heartbeat, each instance reads its little brother's heartbeat and compares it to the value read at the last clock tick. If the new value is greater, or the little brother's ID has changed, the little brother is considered active. However, if the little brother ID and its heartbeat value are the same, the little brother is considered inactive, and the current instance begins watching its little brother's little brother as well. This accumulation of responsibility continues to a predetermined maximum and insures that the failure of one instance does not result in missing the failure of its little brother. If the little brother begins incrementing its heartbeat again, all additional responsibilities are dropped.

If a member instance is judged dead, or disinterested, and it has not notified the APMP computer system of its intent to shut down or crash, the instance is removed from the APMP system. This may be done, for example, by setting the "bugcheck" bit in the instance primitive interrupt mask and sending an IP interrupt to all CPU's of the instance. As a rule, shared memory may only be accessed below the hardware priority of the IP interrupt. This insures that if the CPUs in the instance should attempt to execute at a priority below that of the IP interrupt, the IP interrupt will occur first and thus the CPU will see the "bugcheck" bit before any lower priority threads can execute. This insures the operating system instance will crash and not touch shared resources such as memory which may have been reallocated for other purposes when the instances were judged dead. As an additional or alternative mechanism, a console callback (should one exist) can be invoked to remove the instance. In addition, in accordance with a preferred embodiment, whenever an instance disappears or drops out of the APMP computer system without warning, the remaining instances perform some sanity checks to determine whether they can continue. These checks include verifying that all pages in the APMP database are still accessible, i.e. that there was not a memory failure.

#### Assignment of Resources After Joining

A CPU can have at most one owner partition at any given time in the power-up life of an APMP system. However, the reflection of that ownership and the entity responsible for controlling it can change as a result of configuration and state transitions undergone by the resource itself, the partition it resides within, and the instance running in that partition.

CPU ownership is indicated in a number of ways, in a number of structures dictated by the entity that is managing the resource at the time. In the most basic case, the CPU can be in an unassigned state, available to all partitions that reside in the same sharing set as the CPU. Eventually that CPU is assigned to a specific partition, which may or may not be running an operating system instance. In either case, the partition reflects its ownership to all other partitions through the configuration tree structure, and to all operating system instances that may run in that partition through the AVAILABLE bit in the HWRPB per-CPU flags field.

If the owning partition has no operating system instance running on it, its console is responsible for responding to, and initiating, transition events on the resources within it. The console decides if the resource is in a state that allows it to migrate to another partition or to revert back to the unassigned state.

If, however, there is an instance currently running in the partition, the console relinquishes responsibility for initiating resource transitions and is responsible for notifying the running primary of the instance when a configuration change has taken place. It is still the facilitator of the underlying hardware transition, but control of resource transitions is elevated one level up to the operating system instance. The transfer of responsibility takes place when the primary CPU executes its first instruction outside of console mode in a system boot.

Operating system instances can maintain ownership state information in any number of ways that promote the most efficient usage of the information internally. For example, a hierarchy of state bit vectors can be used which reflect the instance-specific information both internally and globally (to other members sharing an APMP database).

The internal representations are strictly for the use of the instance. They are built up at boot time from the underlying configuration tree and HWRPB information, but are maintained as strict software constructs for the life of the operating system instance. They represent the software view of the partition resources available to the instance, and may--through software rule sets--further restrict the configuration to a subset of that indicated by the physical constructs. Nevertheless, all resources in the partition are owned and managed by the instance - using the console mechanisms to direct state transitions--until that operating system invocation is no longer a viable entity. That state is indicated by halting the primary CPU once again back into console mode with no possibility of returning without a reboot.

Ownership of CPU resources never extends beyond the instance. The state information of each individual instance is duplicated in an APMP database for read-only decision-making purposes, but no other instance can force a state transition event for another's CPU resource. Each instance is responsible for understanding and controlling its own resource set; it may receive external requests for its resources, but only it can make the decision to allow the resources to be transferred.

When each such CPU becomes operational, it does not set its AVAILABLE bit in the per-CPU flags. When the AVAILABLE bit is not set, no instance will attempt to start, nor expect the CPU to join in SMP operation. Instead, the CPU, in console mode, polls the owner field in the configuration tree waiting for a valid partition to be assigned. Once a valid partition is assigned as the owner by the primary console, the CPU will begin operation in that partition.

During runtime, the `current_owner` field reflects the partition where a CPU is executing. The AVAILABLE bit in the per-CPU flags field in the HWRPB remains the ultimate indicator of whether a CPU is actually available, or executing, for SMP operation with an operating system instance, and has the same meaning as in conventional SMP systems.

It should be noted that an instance need not be a member of a sharing set to participate in many of the reconfiguration features of an APMP computer system. An instance can transfer its resources to another instance in the APMP system so that an instance which is not a part of a sharing set can transfer a resource to an instance which is part of the sharing set. Similarly, the instance which is not a part of the sharing set can receive a resource from an instance which is part of the sharing set.

#### Runtime Migration of Resources

After an APMP system is running, resources initially allocated to one partition can move, or migrate, to another partition. This migration may take place under

control of a system administrator or may be initiated by an operating instance without system administrator participation. Migration is accomplished by causing the owning operating system instance to stop using a resource or by a console program taking action to stop using a resource. The configuration tree is then altered to transfer ownership of the resource to another operating system instance. The new operating system instance then begins using the resource. The action required to cause an operating system instance to stop using a resource is operating system specific.

In general, any resource can migrate unless migration is prevented by additional system constraints. For example, any CPU in the APMP system can be moved from a first partition to a second partition, provided it is not currently a primary CPU in the first partition, and it is not bound by operating system constraints, such as distributed interrupt handling. The policy on when and where a CPU may migrate is strictly up to the operating system code which the CPU is executing.

**Memory may also be migrated from a first partition to a second partition. If the memory is privately owned by the first partition, it can be migrated in a straightforward manner. If the memory is shared between two partitions, some additional steps may be necessary to insure that the memory is completely unloaded by all partitions prior to migration.**

Preferably, migration is accomplished without rebooting the entire APMP system, although the migration of certain resources may require a reboot of one or both partitions which are participating in the resource migration depending on the type of change and the operating system capabilities. For example, the migration of an input/output processor from one partition to another may require both partitions to reboot. However, changing the memory allocation of one or more memory fragments may be handled by the operating system code without the need for a reboot.

**The migration process operates under a "push" model in that the original resource owner must first release the resource before the new owner can begin using the resource. The basic steps in one embodiment of this push migration process are illustrated in the flowchart shown in FIG. 9 and schematically in FIGS. 10A-10E.**

In FIG. 9, a resource migration operation commences in step 900 and proceeds to step 902 where an operating system instance (for example, operating instance 2) which is in need of a resource requests use of the resource from another operating system, for example, operating in system instance 1. Interprocessor communication can be carried out by a number of conventional and well-known mechanisms such as interprocess interrupts, common memory or either mechanisms. The same operation step is also illustrated in FIG. 10A. FIG. 10A shows an illustrative system consisting of two partitions: partition 1 (1000) and partition 2 (1002). Each partition 1000 and 1002 includes a console program schematically illustrated as boxes 1006 and 1016, respectively. Each partition 1000 and 1002 also contains an operating system instance illustrated in operating system instance 1 (1008) and operating system instance 2 (1018). Partition 1 (1000) also includes a resource 1010 which is being used by operating system instance 1 (1008) as indicated schematically by arrow 1011. In the first step of the migration process, operating system instance 2 (1018) makes a request, schematically indicated by arrow 1014, to operating system instance 1 (1008) to use resource 1010.

In accordance with the push model used in the present invention, operating system instance 1 (1008) must agree to the transfer of the requested resource 1010 to operating system instance 2 (1018). If operating system instance 1 (1008) agrees, it quiesces the resource 1010 as indicated in step 904. The next step in the process is also illustrated schematically in FIG. 10B in which the elements common to FIG. 10A are indicated with the same numerical designations. For example, partition 1 is labeled 1000 in both FIGS. 10A and 10B. As shown in

FIG. 10B, If operating system instance 1 (1008) agrees to migrate the resource, operating system instance 1, 1008 quiesces or stops its use of resource 1010 as indicated schematically by dotted arrow 1020.

Next, operating system instance 1 (1008) informs console 1 (1006) of the intended change of ownership of the resource as indicated in step 906. This step in the migration process is illustrated in FIG. 10C. In this step, operating system instance 1 (1008) informs console 1 (1006) that a transfer of ownership is requested as indicated schematically by arrow 1022. In particular, operating system instance 1 (1008) informs console 1 (1006) to change the owner and current\_owner fields for resource 1010 in configuration tree 1012, as described below.

Next, in step 908, console 1 (1006) modifies the configuration tree 1012 to indicate the new owner of the resource. This step is illustrated schematically in FIG. 10D where console 1 (1006) modifies configuration tree 1012 as indicated schematically by arrow 1026. This modification effectively moves resource 1010 from partition 1 (1000) to partition 2 (1002) as shown schematically by arrow 1028 so that the resource now appears in partition 2 (1002) as shown schematically by element 1024.

Finally, in step 910, operating system instance 1 (1008) in conjunction with console 1 (1006) informs operating system instance 2 (1018) in conjunction with console 2 (1016) of the transfer so that operating system instance 2 (1018) can begin using the resource. The routine then ends in step 912. This final step in the process is illustrated by FIG. 10E in which operating system instance 1 (1008) informs operating system instance 2 (1018) that the resource 1024 is now available. This operation is shown schematically by arrow 1030 and involves known and conventional interprocess communications. Operating system instance 2 (1018) may then use resource 1024 as shown schematically by arrow 1032.

As previously mentioned, owner and current\_owner fields are provided in each node of the configuration tree to manage the transition of a resource from an owned and active state, to a un-owned and inactive state. Only the owning operating system instance of a resource can de-assign the resource by setting the owner field to zero. Only resources which have null owner and current\_owner fields may be assigned to a new partition or community. During a resource migration, the HWRPBs and configuration trees are modified in concert and atomically to prevent any instance from obtaining an incorrect view of the configuration. In a similar fashion, a "hot swap" must also be accompanied by atomic and coordinated modifications to the HWRPBs and configuration trees. When a resource is de-assigned, the owning operating system may choose to de-assign the owner field, or both the owner and current\_owner fields. The decision of which field to de-assign is based on the ability of the owning operating system to discontinue the use of the resource prior to de-assigning ownership. In the case where a reboot is required to relinquish ownership, the owner field is cleared, but the current\_owner field is not changed. When the owning operating system instance reboots, the console program in the partition can be designed to clear any current\_owner fields for resources that have null owner fields during initialization.

In the case where the resource is shared among operating system instances which are part of a community, the operating system instances must cooperate in order to de-assign the resource. This de-assignment is managed by the instances which are part of the community.

Ownership of a resource is altered by changing the configuration tree. Certain rules must be followed when the configuration tree is altered. They include the following:

- 1) If the config field of a configuration tree node points to the tree hardware root node, the corresponding resource is capable of being assigned independently to any community or partition.

2) If the config field of that node does not point to the tree hardware root node, when the resource is assigned to a partition, all descendants of the corresponding tree node must be modified so that they are assigned to the same partition and all tree nodes in the parent chain from the tree node being altered to the config node must also have the same Partition, or Community owner. The config node must be an ancestor of the node being changed and reachable by following the parent pointers.

When the ownership of a hardware node is given to a Community, all descendants must either be owned by the community, a partition in the community, or unowned. If the config pointer is not the hardware root, then all nodes in the parent chain from the node being altered to the config node must also have the same Community owner. The config node must be an ancestor of the node being changed and reachable by following the parent pointers.

Hardware components, for example a CPU, may be free to operate independently. In this case, the config pointer specifies the hardware root. However, components may also be made up of other components that cannot operate independently. For example, the I/O Controllers on a PCI bus might not be separable from the PCI bus.

Some hardware components, such as memory may be shared by multiple partitions. This is expressed by the owner field pointing to a community. Partitions in the community, may then share access to the component. The descendants of the hardware component that is shared may specify an owner that is a descendent. For example, a memory subsystem may be owned by the community. In this case, the memory controller, which expresses the hardware aspects of the memory, including error handling, is owned by a partition in the community, and the memory descriptor node may be owned by the community, and its fragments owned by both communities (shared) and by partitions (private).

A software implementation of the above-described embodiment may comprise a series of computer instructions either fixed on a tangible medium, such as a computer readable media, e.g. a diskette, a CD-ROM, a ROM, or fixed disk, or transmittable to a computer system, via a modem or other interface device, such as a communications adapter connected to the network over a medium. The medium can be either a tangible medium, including but not limited to optical or analog communications lines, or may be implemented with wireless techniques, including but not limited to microwave, infrared or other transmission techniques. It may also be the Internet. The series of computer instructions embodies all or part of the functionality previously described herein with respect to the invention. Those skilled in the art will appreciate that such computer instructions can be written in a number of programming languages for use with many computer architectures or operating systems. Further, such instructions may be stored using any memory technology, present or future, including, but not limited to, semiconductor, magnetic, optical or other memory devices, or transmitted using any communications technology, present or future, including but not limited to optical, infrared, microwave, or other transmission technologies. It is contemplated that such a computer program product may be distributed as a removable media with accompanying printed or electronic documentation, e.g., shrink wrapped software, pre-loaded with a computer system, e.g., on system ROM or fixed disk, or distributed from a server or electronic bulletin board over a network, e.g., the Internet or World Wide Web.

Although an exemplary embodiment of the invention has been disclosed, it will be apparent to those skilled in the art that various changes and modifications can be made which will achieve some of the advantages of the invention without departing from the spirit and scope of the invention. For example, it will be obvious to those reasonably skilled in the art that, although the description was directed to a particular hardware system and operating system, other hardware and operating system software could be used in the same manner as that described. Further, other migration scenarios are possible. For example, a



system manager could direct the first operating system instance to quiesce a resource and then place the resource into the "unassigned" state. Some time later, a second instance could discover a need for the resource and remove it from the pool of unassigned resources. Alternatively, the system manager could direct the second instance to allocate the resource. Alternatively, system policy might cause a batch job to run at a predetermined time on the first instance that would unilaterally quiesce and transfer some resource to a second instance. In both of these latter examples, there is no explicit communication between the instances regarding the resource, although there is clearly some human intervention or policy that coordinates the overall usage of resources. In yet another approach, instances would allocate a resource from the unassigned pool when their needs require it. When they are finished with the resource, they return it to the unassigned pool. If the pool is empty, an instance has to operate with its then-existing allocation of the resource. So, we have a common pool, but no request/release protocol is required between instances.

Other aspects, such as the specific instructions utilized to achieve a particular function, as well as other modifications to the inventive concept are intended to be covered by the appended claims.

What is claimed is:

1. A computer system having a plurality of assignable system resources, including processors, memory and I/O circuitry, the computer system comprising: an interconnection mechanism for electrically interconnecting the processors, the memory and the I/O circuitry so that each processor has electrical access to all of the memory and at least some of the I/O circuitry; a software mechanism for assigning the assignable system resources to a plurality of partitions, each partition including at least one processor, a portion of the memory and a portion of the I/O circuitry; an operating system instance running in each partition; and a migration mechanism for migrating one or more of the assignable system resources from one partition to another partition.
2. A computer system according to claim 1, wherein the migration mechanism comprises a runtime migration mechanism for moving the one or more of the assignable system resources during system operation without a reboot of the entire system.
3. A computer system according to claim 1, wherein the one or more of the assignable system resources includes a first processor, wherein the migration mechanism comprises a first communication mechanism which communicates with an operating system instance in a partition from which the first processor is being moved to cause the operating system instance to quiesce the first processor.
4. A computer system according to claim 3, wherein an operating system instance in a partition from which the first processor is being moved must acquiesce to the move before the first processor can be removed.
5. A computer system according to claim 1 wherein the software mechanism for assigning system resources comprises a first data storage for maintaining configuration information indicating which of the plurality of processors is assigned to each partition.
6. A computer system according to claim 5, wherein the software mechanism for assigning resources comprises a mechanism for modifying the configuration information to change an assignment of assignable system resources from one partition to another partition.
7. A computer system according to claim 1, wherein the one or more of the assignable system resources includes a second processor, wherein the migration mechanism comprises a second communication mechanism which communicates with an operating system instance in a partition to which the second processor is being moved to cause the operating system instance to begin using the second processor.

8. A computer system according to claim 1, wherein the one or more of the assignable system resources includes a first region of the memory, wherein the migration mechanism comprises a third communication mechanism which communicates with an operating system instance in a partition from which the first region of the memory is being moved to cause the operating system instance to unload the first region of the memory.
9. A computer system according to claim 1, wherein the one or more of the assignable system resources includes a second region of the memory, wherein the migration mechanism comprises a fourth communication mechanism which communicates with a plurality of operating system instances which share the second region of the memory to be moved to cause the plurality of the operating system instances to cooperatively unload the second region of the memory.
10. A computer system according to claim 9, wherein the migration mechanism comprises a console controlled by one program of the plurality of operating system instances which share the second region of the memory to move the second region of the memory from the one partition to the other partition.
11. A computer system according to claim 1, wherein the software mechanism for assigning resources comprises a second data storage which maintains configuration information indicating which portions of the memory are assigned to each partition.
12. A computer system according to claim 1, wherein the one or more of the assignable system resources includes a third region of the memory, wherein the migration mechanism comprises a fifth communication mechanism which communicates with an operating system instance in a partition to which the third region of the memory is being moved to cause the operating system instance to begin filling the third region of the memory.
13. A computer system according to claim 1, wherein the plurality of processors is physically divided into partitions and wherein each respective partition comprises a console program which interacts with the processors in the respective partition.
14. A computer system according to claim 13, wherein the software mechanism for assigning resources comprises a console program which maintains configuration information indicating which of the plurality of processors, which portions of the memory, and which portions of the I/O circuitry are assigned to each partition.
15. A method for changing a configuration of a computer system, the method comprising the steps of: (a) electrically interconnecting a plurality of assignable system resources, including processors, memory and I/O circuitry, so that each processor has electrical access to all of the memory and at least some of the I/O circuitry; (b) assigning the assignable system resources to a plurality of partitions, each partition including at least one processor, a portion of the memory, and a portion of the I/O circuitry; (c) running an operating system instance in each partition; and (d) changing the configuration by migrating one or more of the assignable system resources from one partition to another partition.
16. A method according to claim 15, wherein step (d) comprises the step of: (d1) moving one or more of the assignable system resources from one partition to another partition during system operation without a reboot of the entire system.
17. A method according to claim 15 wherein step (d) comprises the step of: (d2) communicating with an operating system instance in a partition from which a processor is being moved to cause the operating system instance to quiesce the processor.
18. A method according to claim 17, wherein step (d2) comprises the step of: (d2a) requiring an operating system instance in the partition from which the processor is being moved to acquiesce to the move before the processor can be removed from the partition.

19. A method according to claim 15 wherein step (b) comprise the step of: (b1) maintaining, in a first data storage, configuration information indicating which of the plurality of processors is assigned to each partition.
20. A method according to claim 19, wherein step (b) comprises the step of: (b2) modifying the configuration information to change an assignment of one or more of the assignable system resources from one partition to another partition.
21. A method according to claim 15 wherein step (d) comprises the step of: (d3) communicating with an operating system instance in a partition to which a processor is being moved to cause the operating system instance to begin using the processor.
22. A method according to claim 15, wherein step (d) comprises the step of: (d4) communicating with an operating system instance in a partition from which a region of the memory is being moved to cause the operating system instance to unload the region of the memory.
23. A method according to claim 15, wherein step (d) comprises the step of: (d5) communicating with a plurality of operating system instances which share a region of the memory to be moved to cause the plurality of operating system instances to cooperatively unload the region of the memory to be moved.
24. A method according to claim 23, wherein step (d5) comprises the step of: (d5a) controlling a console program with one of the plurality of operating system instances to move the region of the memory from the one partition to the other partition.
25. A method according to claim 15, wherein step (b) comprises the step of: (b3) maintaining, in a second data storage, configuration information indicating which portions of the memory are assigned to each partition.
26. A method according to claim 15, wherein step (d) comprises the step of: (d6) communicating with an operating system instance in a partition to which a region of the memory is being moved to cause the operating system instance to begin filling the region of the memory.
27. A method according to claim 15 wherein the plurality of processors is physically divided into partitions and wherein the method further comprises the step of: (e) providing each partition with a console program which interacts with the processors in the partition.
28. A method according to claim 27 wherein step (e) comprises the step of: (e1) providing a console program which maintains configuration information indicating which of the plurality of processors, which memory portions and which I/O circuitry are assigned to each partition.
29. A computer program product for changing a configuration of a computer system having a plurality of assignable system resources, including processors, memory and I/O circuitry, in which the processors, memory, and I/O circuitry are electrically interconnected so that each processor has electrical access to all of the memory and at least some of the I/O circuitry, the computer program product comprising a computer usable medium having computer readable program code thereon, comprising: (a) program code which assigns the assignable system resources to a plurality of partitions, each partition including at least one processor, a portion of the memory, and a portion of the I/O circuitry; (b) operating system code for creating an operating system instance running in each partition; and (c) program code which changes the configuration by migrating one or more of the assignable system resources from one partition to another partition.
30. A computer program product according to claim 29, wherein the program code which changes the configuration comprises program code which moves one or more of the assignable system resources from one partition to another partition during system operation without a reboot of the entire system.
31. A computer program product according to claim 29 wherein the program code which changes the configuration comprises program code which communicates with

an operating system instance in a partition from which a processor is being moved to cause the operating system instance to quiesce the processor.

32. A computer product according to claim 31, wherein the program code which communicates with an operating system instance comprises program code which requires the operating system instance in the partition from which the processor is being moved to acquiesce to the move before the processor can be removed.

33. A computer program product according to claim 29 wherein the program code which assigns the system resources comprises program code which maintains, in a first data storage, configuration information indicating which of the plurality of processors is assigned to each partition.

34. A computer program product according to claim 33, wherein the program code which assigns the system resources comprises program code which modifies the configuration information to change the assignment of one or more of the assignable system resources from one partition to another partition.

35. A computer program product according to claim 29 wherein the program code which changes the configuration comprises program code which communicates with an operating system instance in a partition to which a processor is being moved to cause the operating system instance to begin using the processor.

36. A computer program product according to claim 29, wherein the program code which changes the configuration comprises program code which communicates with an operating system instance in a partition from which a region of the memory is being moved to cause the operating system instance to unload the region of the memory being moved.

37. A computer program product according to claim 29, wherein the program code which changes the configuration comprises program code which communicates with a plurality of operating system instances which share a region of the memory to be moved to cause the plurality of operating system instances to cooperatively unload the region of the memory to be moved.

38. A computer program product according to claim 37, wherein the program code which communicates with a plurality of operating system instances comprises program code in one of the plurality of operating system instances which controls a console program to move the region of the memory from the one partition to the other partition.

39. A computer program product according to claim 29, wherein the program code which assigns the system resources comprises program code which maintains, in a second data storage, configuration information indicating which portions of the memory are assigned to each partition.

40. A computer program product according to claim 29, wherein the program code which changes the configuration comprises program code which communicates with an operating system instance in a partition to which a region of the memory is being moved to cause the operating system instance to begin filling the region of the memory.

41. A computer program product according to claim 29 wherein the plurality of processors is physically divided into partitions and wherein the computer program product further comprises console program code which interacts with the processors in the partition.

42. A computer program product according to claim 41, wherein the console program code comprises console program code which maintains configuration information indicating which of the plurality of processors, which portions of the memory and which portions of the I/O circuitry are assigned to each partition.

ISSUE U.S. PATENT CLASSIF.:

MAIN: 709/226.000

SECONDARY: 707/010.000; 707/202.000; 709/104.000; 712/028.000;

714/002.000

CURRENT U.S. PATENT CLASSIF.:

MAIN: 709/226.000  
SECONDARY: 707/010.000; 707/202.000; 709/104.000; 712/028.000;  
714/002.000  
INT. PATENT CLASSIF.: [7]  
MAIN: G06F009-40  
FIELD OF SEARCH: 707/10; 707/100; 707/202; 707/223; 709/104; 709/106;  
709/226; 712/28; 712/203; 713/2; 713/100; 714/2; 714/4  
ART UNIT: 212